

COMPUTATIONAL COMPLEXITY VERSUS STATISTICAL PERFORMANCE ON SPARSE RECOVERY PROBLEMS

VINCENT ROULET, NICOLAS BOUMAL, AND ALEXANDRE D'ASPREMONT

ABSTRACT. We show that several classical quantities controlling compressed sensing performance directly match classical parameters controlling algorithmic complexity. We first describe linearly convergent restart schemes on first-order methods solving a broad range of compressed sensing problems, where sharpness at the optimum controls convergence speed. We show that for sparse recovery problems, this sharpness can be written as a condition number, given by the ratio between true signal sparsity and the largest signal size that can be recovered by the observation matrix. In a similar vein, Renegar's condition number is a data-driven complexity measure for convex programs, generalizing classical condition numbers for linear systems. We show that for a broad class of compressed sensing problems, the worst case value of this algorithmic complexity measure taken over all signals matches the restricted singular value of the observation matrix which controls robust recovery performance. Overall, this means in both cases that, in compressed sensing problems, a single parameter directly controls both computational complexity and recovery performance. Numerical experiments illustrate these points using several classical algorithms.

INTRODUCTION

Sparse recovery problems have received a lot of attention from various perspectives. On one side, an extensive literature explores the limits of recovery performance. On the other side, a long list of algorithms now solve these problems very efficiently. Early on, it was noticed empirically by e.g. Donoho and Tsaig [2008], that recovery problems which are easier to solve from a statistical point of view (i.e., where more samples are available), are also easier to solve numerically. Here, we show that these two aspects are indeed intimately related.

Recovery problems consist in retrieving a signal x^* , lying in some Euclidean space E , given linear observations. If the signal is “sparse”, namely if it can be efficiently compressed, a common approach is to minimize the corresponding sparsity inducing norm $\|\cdot\|$ (e.g. the ℓ_1 norm in classical sparse recovery). The exact sparse recovery problem then reads

$$\begin{aligned} & \text{minimize} && \|x\| \\ & \text{subject to} && A(x) = b, \end{aligned} \tag{1}$$

in the variable $x \in E$, where A is a linear operator on E and $b = A(x^*)$ is the vector of observations. If the observations are affected by noise a robust version of this problem is written as

$$\begin{aligned} & \text{minimize} && \|x\| \\ & \text{subject to} && \|A(x) - b\|_2 \leq \epsilon, \end{aligned}$$

in the variable $x \in E$, where $\|\cdot\|_2$ is the Euclidean norm and $\epsilon > 0$ is a tolerance to noise. In penalized form, this is

$$\text{minimize} \quad \|x\| + \lambda \|A(x) - b\|_2^2$$

in the variable $x \in E$ where $\lambda > 0$ is a penalization parameter. This last problem is known as the LASSO [Tibshirani, 1996] in the ℓ_1 case.

Date: November 7, 2017.

2010 Mathematics Subject Classification. 90C25, 94A12.

Key words and phrases. Renegar's condition number, distance to infeasibility, sharpness, restart, error bounds, sparse recovery.

When x^* has no more than s non zero values, [Donoho and Tanner \[2005\]](#) and [Candès and Tao \[2006\]](#) have shown that, for certain linear operators A , $O(s \log p)$ observations suffice for stable recovery of x^* by solving the exact formulation (1) using the ℓ_1 norm (a linear program). These results have then been generalized to many other recovery problems with various assumptions on signal structure (e.g., where x is a block-sparse vector, a low-rank matrix, etc.) and corresponding convex relaxations were developed in those cases (see e.g. [Chandrasekaran et al. \[2012\]](#) and references therein). Recovery performance is often measured in terms of the number of samples required to guarantee exact or robust recovery given a level of noise.

On the computational side, many algorithms were developed to solve these problems at scale. Besides specialized methods such as LARS [[Efron et al., 2004](#)], FISTA [[Beck and Teboulle, 2009](#)] and NESTA [[Becker, Bobin and Candès, 2011](#)], solvers use accelerated gradient methods to solve robust recovery problems, with efficient and flexible implementations covering a wide range of compressed sensing instances developed by e.g. [Becker, Candès and Grant \[2011\]](#). Recently, linear convergence results have been obtained for the LASSO [[Agarwal et al., 2011](#); [Yen et al., 2014](#); [Zhou et al., 2015](#)] using variants of the classical strong convexity assumption, while [[Zhou and So, 2017](#)] studied error bounds for a much broader class of structured optimization problems including sparse recovery and matrix completion. Some restart schemes have also been developed in e.g. [[O’Donoghue and Candès, 2015](#); [Su et al., 2014](#); [Giselsson and Boyd, 2014](#)] while [Fercoq and Qu \[2016\]](#) showed that generic restart schemes can offer linear convergence given a rough estimate of the behavior of the function around its minimizers.

As mentioned above, [Donoho and Tsaig \[2008\]](#) was one of the first reference to connect statistical and computational performance in this case, showing empirically that recovery problems which are easier to solve from a statistical point of view (i.e., where more samples are available), are also easier to solve numerically (using homotopy methods). More recently, [Chandrasekaran and Jordan \[2013\]](#); [Amelunxen et al. \[2014\]](#) studied computational and statistical tradeoffs for increasingly tight convex relaxations of shrinkage estimators. They show that recovery performance is directly linked to the Gaussian squared-complexity of the tangent cone with respect to the constraint set and study the complexity of several convex relaxations. In [[Chandrasekaran and Jordan, 2013](#); [Amelunxen et al., 2014](#)] however, the structure of the convex relaxation is varying and affecting both complexity and recovery performance, while in [[Donoho and Tsaig, 2008](#)] and in what follows, the structure of the relaxation is fixed, but the data (i.e. the observation matrix A) varies.

Here, as a first step, we study the exact recovery case and show that the null space property introduced by [Cohen et al. \[2009\]](#) can be seen as a measure of sharpness on the optimum of the sparse recovery problem. On one hand this allows us to develop linearly convergent restart schemes whose rate depends on this sharpness. On the other hand we recall how the null space property is linked to the recovery threshold of the sensing operator A for random designs, thus producing a clear link between statistical and computational performance.

We then analyze the underlying conic geometry of recovery problems. Robust recovery performance is controlled by a minimal conically restricted singular value. We recall Renegar’s condition number and show how it affects the computational complexity of optimality certificates for exact recovery and the linear convergence rate of restart schemes. By observing that the minimal conically restricted singular value matches the worst case value of Renegar’s condition number on sparse signals, we provide further evidence that a single quantity controls both computational and statistical aspects of recovery problems. Numerical experiments illustrate its impact on various classical algorithms for sparse recovery.

The first two sections focus on the ℓ_1 case for simplicity. We generalize our results to non-overlapping group norms and the nuclear norm in a third section.

Notations. For a given integer $p \geq 1$, $\llbracket 1, p \rrbracket$ denotes the set of integers between 1 and p . For a given subset $S \subset \llbracket 1, p \rrbracket$, we denote $S^c = \llbracket 1, p \rrbracket \setminus S$ its complementary and $\mathbf{Card}(S)$ its cardinality. For a given vector $x \in \mathbb{R}^p$, we denote $\text{Supp}(x) = \{i \in \llbracket 1, p \rrbracket : x_i \neq 0\}$ the support of x , $\|x\|_0 = \mathbf{Card}(\text{Supp}(x))$ its sparsity and $\|x\|_p$ its p -norm. For a given vector x and integer subset $S \subset \llbracket 1, p \rrbracket$, $x_S \in \mathbb{R}^p$ denotes the vector

obtained by zeroing all coefficients of x that are not in S . For a given linear operator or matrix A , we denote $\text{Null}(A)$ its null space, $\text{Im}(A)$ its range, and $\|X\|_2$ its operator norm with respect to the Euclidean norm (for matrices this is the spectral norm). The identity operator is denoted \mathbf{I} . In a linear topological space E we denote $\text{int}(F)$ the interior of $F \subset E$. Finally for a given real a , we denote $\lceil a \rceil$ the smallest integer larger than or equal to a and $\lfloor a \rfloor$ the largest integer smaller than or equal to a .

1. SHARPNESS, RESTART AND SPARSE RECOVERY PERFORMANCE

In this section and the following one, we discuss sparse recovery problems using the ℓ_1 norm. Given a matrix $A \in \mathbb{R}^{n \times p}$ and observations $b = Ax^*$ on a signal $x^* \in \mathbb{R}^p$, recovery is performed by solving the ℓ_1 minimization program

$$\begin{aligned} & \text{minimize} && \|x\|_1 \\ & \text{subject to} && Ax = b \end{aligned} \tag{\ell_1 \text{ recovery}}$$

in the variable $x \in \mathbb{R}^p$.

In what follows, we show that the Null Space Property condition (recalled below) can be seen as measure of sharpness for ℓ_1 -recovery of a sparse signals x^* , with

$$\|x\|_1 - \|x^*\|_1 > \gamma \|x - x^*\|_1 \tag{Sharp}$$

for any $x \neq x^*$ such that $Ax = b$, and some $0 \leq \gamma < 1$. This first ensures that x^* is the unique minimizer of problem (ℓ_1 recovery) but also has important computational implications. It allows us to produce linear convergent restart schemes whose rates depend on sharpness. By connecting the null space property to recovery threshold for random observation matrices, we get a direct link between computational and statistical aspects of sparse recovery problems.

1.1. Null space property & sharpness for exact recovery. Although the definition of null space property appeared in earlier work [Donoho and Huo, 2001; Feuer and Nemirovski, 2003] the terminology of restricted null space is due to Cohen et al. [2009]. The following definition differs slightly from the original one in order to relate it to intrinsic geometric properties of the problem in Section 2.

Definition 1.1. (Null Space Property) *The matrix A satisfies the Null Space Property (NSP) on support $S \subset \llbracket 1, p \rrbracket$ with constant $\alpha \geq 1$ if for any $z \in \text{Null}(A) \setminus \{0\}$,*

$$\alpha \|z_S\|_1 < \|z_{S^c}\|_1. \tag{NSP}$$

The matrix A satisfies the Null Space Property at order s with constant $\alpha \geq 1$ if it satisfies it on every support S of cardinality at most s .

The Null Space Property is a necessary and sufficient condition for the convex program (ℓ_1 recovery) to recover all signals up to some sparsity threshold. Necessity will follow from results recalled in Section 2.2. We detail sufficiency of (NSP) using sharpness in the following proposition.

Proposition 1.2. *Given a coding matrix $A \in \mathbb{R}^{n \times p}$ satisfying (NSP) at order s with constant $\alpha \geq 1$, if the original signal x^* is s -sparse, then for any $x \in \mathbb{R}^p$ satisfying $Ax = b$, $x \neq x^*$, we have*

$$\|x\|_1 - \|x^*\|_1 > \frac{\alpha - 1}{\alpha + 1} \|x - x^*\|_1.$$

This implies signal recovery, i.e. optimality of x^ for (ℓ_1 recovery), and the sharpness bound (Sharp) with $\gamma = \frac{\alpha - 1}{\alpha + 1}$.*

Proof. The proof follows the one in [Cohen et al. \[2009, Theorem 4.4\]](#). Let $S = \text{supp}(x^*)$, with $\text{Card}(S) \leq s$, and let $x \neq x^*$ such that $Ax = b$, so $z = x - x^* \neq 0$ satisfies $Az = 0$. Then

$$\begin{aligned} \|x\|_1 &= \|x_S^* + z_S\|_1 + \|z_{S^c}\|_1 \\ &\geq \|x_S^*\|_1 - \|z_S\|_1 + \|z_{S^c}\|_1 \\ &= \|x^*\|_1 + \|z\|_1 - 2\|z_S\|_1. \end{aligned}$$

Now as A satisfies **(NSP)** on support S ,

$$\|z\|_1 = \|z_S\|_1 + \|z_{S^c}\|_1 > (1 + \alpha)\|z_S\|_1$$

hence

$$\|x\|_1 - \|x^*\|_1 > \frac{\alpha - 1}{\alpha + 1} \|z\|_1 = \frac{\alpha - 1}{\alpha + 1} \|x - x^*\|_1.$$

As $\alpha \geq 1$, this implies that x^* is the solution of program (**ℓ_1 recovery**) and the corresponding sharpness bound. ■

Sharpness is a crucial property for optimization problems that can be exploited to accelerate the performance of classical optimization algorithms [[Nemirovskii and Nesterov, 1985](#); [Roulet and d'Aspremont, 2017](#)]. Before that we remark that it is in fact equivalent to **(NSP)** at order s .

Proposition 1.3. *Given a matrix $A \in \mathbb{R}^{n \times p}$ such that problem (**ℓ_1 recovery**) is sharp on every s -sparse signal x^* , i.e. there exists $0 \leq \gamma < 1$ such that*

$$\|x\|_1 - \|x^*\|_1 > \gamma \|x - x^*\|_1,$$

for any $x \neq x^*$ such that $Ax = Ax^*$. Then, A satisfies **(NSP)** at order s with constant $\alpha = \frac{1+\gamma}{1-\gamma} \geq 1$.

Proof. Let $S \subset \llbracket 1, p \rrbracket$ with $\text{Card}(S) \leq s$ and $z \in \text{Null}(A)$, $z \neq 0$, such that $Az_S = -Az_{S^c}$ and $z_S \neq -z_{S^c}$. Using sharpness of problem (**ℓ_1 recovery**) with $x^* = z_S$, and $x = -z_{S^c}$, we get

$$\|z_{S^c}\|_1 - \|z_S\|_1 > \gamma \|z\|_1 = \gamma \|z_S\|_1 + \gamma \|z_{S^c}\|_1.$$

Rearranging terms and using $\gamma < 1$, this reads

$$\|z_{S^c}\|_1 > \frac{1 + \gamma}{1 - \gamma} \|z_S\|_1,$$

which is **(NSP)** on support S with the announced constant. As S was taken arbitrarily, this means **(NSP)** holds at order s . ■

1.2. Restarting first-order methods. In this section, we seek to solve the recovery problem (**ℓ_1 recovery**) and exploit the sharpness bound (**Sharp**). The NESTA algorithm [[Becker, Bobin and Candès, 2011](#)] uses the smoothing argument of [Nesterov \[2005\]](#) to solve (**ℓ_1 recovery**). In practice, this means using the optimal algorithm of [Nesterov \[1983\]](#) to minimize

$$f_\epsilon(x) \triangleq \sup_{\|u\|_\infty \leq 1} \left\{ u^T x - \frac{\epsilon}{2p} \|u\|_2^2 \right\}$$

for some $\epsilon > 0$, which approximates the ℓ_1 norm uniformly up to $\epsilon/2$. This is the classical Huber function, which has a Lipschitz continuous gradient with constant equal to p/ϵ . Overall given an accuracy ϵ and a starting point x_0 this method outputs after t iterations a point $x = \mathcal{A}(x_0, \epsilon, t)$ such that

$$\|x\|_1 - \|\hat{x}\|_1 \leq \frac{2p\|x_0 - \hat{x}\|_2^2}{\epsilon t^2} + \frac{\epsilon}{2}, \quad (2)$$

for any \hat{x} solution of problem (**ℓ_1 recovery**). Now if the sharpness bound is satisfied, restarting this method, as described in the (**Restart**) scheme presented below, accelerates its convergence.

Algorithm 1 Restart Scheme (**Restart**)

Input: Initial point $y_0 \in \mathbb{R}^p$, initial gap $\epsilon_0 \geq \|y_0\|_1 - \|\hat{x}\|_1$, decreasing factor ρ , restart clock t
For $k = 1 \dots, K$ compute

$$\epsilon_k = \rho\epsilon_{k-1}, \quad y_k = \mathcal{A}(y_{k-1}, \epsilon_k, t) \quad (\text{Restart})$$

Output: A point $\hat{y} = y_K$ approximately solving (ℓ_1 recovery).

1.2.1. *Optimal restart scheme.* We begin by analyzing an optimal restart scheme assuming the sharpness constant is known. We use a non-integer clock to highlight its dependency to the sharpness. Naturally clock and number of restarts must be integer but this does not affect much bounds as detailed in Appendix A. The next proposition shows that algorithm \mathcal{A} needs a constant number of iterations to decrease the gap by a constant factor, which means restart leads to linear convergence.

Proposition 1.4. *Given a coding matrix $A \in \mathbb{R}^{n \times p}$ and a signal $x^* \in \mathbb{R}^p$ such that the sharpness bound (**Sharp**) is satisfied with $\gamma > 0$, i.e.*

$$\|x\|_1 - \|x^*\|_1 > \gamma \|x - x^*\|_1,$$

for any $x \neq x^*$ such that $Ax = Ax^*$, running the (**Restart**) scheme with $t \geq \frac{2\sqrt{p}}{\gamma\rho}$ ensures

$$\|y_k\|_1 - \|x^*\|_1 \leq \epsilon_k, \quad (3)$$

at each iteration, with x^* the unique solution of problem (ℓ_1 recovery). Using optimal parameters

$$\rho^* = e^{-1} \quad \text{and} \quad t^* = \frac{2e\sqrt{p}}{\gamma}, \quad (4)$$

we get a point \hat{y} such that

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \exp\left(-\frac{\gamma}{2\sqrt{p}}eN\right)\epsilon_0. \quad (5)$$

after running a total of N inner iterations of Algorithm **Restart** with $t = t^*$ (hence N/t restarts).

Proof. By the choice of ϵ_0 , (3) is satisfied for $k = 0$. Assuming it holds at iteration k , combining sharpness bound (**Sharp**) and complexity bound (2) leads to, for $x = \mathcal{A}(y_{k-1}, \epsilon_k, t)$,

$$\begin{aligned} \|x\|_1 - \|x^*\|_1 &\leq \frac{2p(\|y_{k-1}\|_1 - \|x^*\|_1)^2}{\gamma^2\epsilon_k t^2} + \frac{\epsilon_k}{2} \\ &\leq \frac{4p}{\rho^2\gamma^2 t^2} \frac{\epsilon_k}{2} + \frac{\epsilon_k}{2}. \end{aligned}$$

Therefore after $t \geq \frac{2\sqrt{p}}{\gamma\rho}$ iterations, the method has achieved the decreased accuracy ϵ_k which proves (3). The overall complexity after a total of N inner iterations, hence N/t restarts, is then

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \rho^{N/t}\epsilon_0.$$

If γ is known, using exactly $\frac{2\sqrt{p}}{\gamma\rho}$ inner iterations at each restart leads to

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \exp\left(\frac{\gamma}{2\sqrt{p}}N\rho \log \rho\right)\epsilon_0.$$

Optimizing in ρ yields $\rho^* = e^{-1}$, and with t^* inner iterations the complexity bound (5) follows. ■

To run NESTA, $A^T A$ is assumed to be an orthogonal projector (often the case in compressed sensing applications) such that the projection on the feasible set is easy. Becker, Bobin and Candès [2011] already studied restart schemes that they called ‘‘acceleration with continuation’’. However their restart criterion depends on the relative variation of objective values, not on the number of iterates, and no linear convergence

was proven. We further note that linear convergence of restart schemes only requires an assumption of the form

$$f(x) - f^* \geq \gamma d(x, X^*), \quad (6)$$

where $d(x, X^*)$ is the distance (in any norm) from x to the set of minimizers of the objective function f (here $f(x) = \|x\|_1$). This type of bound is known as Łojasiewicz's inequality, studied for example in Bolte et al. [2007] for non-smooth convex functions. Here (NSP) ensures that the set of minimizers is reduced to a singleton, the original signal.

1.2.2. Practical restart scheme. Several parameters are needed to run the optimal scheme above. The optimal decreasing factor is independent of the data. The initial gap ϵ_0 can be taken as $\|y_0\|_1$ for $Ay_0 = b$. The sharpness constant γ is for its part mostly unknown such that we cannot choose the number t^* of inner iterations a priori. However, given a budget of iterations N , a log scale grid search can be performed on the optimal restart clock to get nearly optimal rates as detailed in the following corollary (contrary to the general results in [Roulet and d'Aspremont, 2017], the sharpness exponent is known here, simplifying the parameter search).

Corollary 1.5. *Given a coding matrix $A \in \mathbb{R}^{n \times p}$, a signal $x^* \in \mathbb{R}^p$ such that the sharpness bound (Sharp) is satisfied with $\gamma > 0$, a budget of N iterations, run the following schemes from an initial point y_0*

$$\text{(Restart) with } t = h^j, \quad j = 1, \dots, \lfloor \log_h N \rfloor$$

with h the grid search precision. Stop restart iteration when the total number of iterations has exceeded the budget N . Then, provided that $N \geq ht^$, where t^* is defined in (4), at least one of these restart schemes achieves a precision given by*

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \exp\left(-\frac{\gamma}{2h\sqrt{p}}eN\right)\epsilon_0. \quad (7)$$

Overall running the logarithmic grid search has a complexity $\log_h N$ times higher than running N iterations in the optimal scheme.

Proof. All schemes stop after at most $N + h^j \leq 2N$ iterations. As we assumed $N \geq ht^*$, $j = \lfloor \log_h t^* \rfloor \leq \log_h N$ and (Restart) has been run with $t = h^j$. Proposition 1.4 ensures, since $t \geq t^*$, that the output of this scheme achieves after $N' \geq N$ total iterations a precision

$$\|\hat{y}\|_1 - \|\hat{x}\|_1 \leq e^{-N'/t}\epsilon_0 \leq e^{-N/t}\epsilon_0$$

and as $t \leq ht^*$

$$\|\hat{y}\|_1 - \|\hat{x}\|_1 \leq e^{-N/(ht^*)}\epsilon_0$$

which gives the result. Finally the logarithmic grid search costs $\log_h N$ to get this approximative optimal bound. ■

Sharpness therefore controls linear convergence of simple restart schemes to solve (ℓ_1 recovery). We now turn back to (NSP) estimates and connect them to recovery thresholds of the sampling matrix. This will give us a direct link between computational complexity and recovery performance for exact recovery problems.

1.3. Recovery threshold. If (NSP) is satisfied at a given order s it holds also for any $s' \leq s$. However, the constant, and therefore the speed of convergence, may change. Here we show that this constant actually depends on the ratio between the maximal order at which A satisfies (NSP) and the sparsity of the signal that we seek to recover.

To this end, we give a more concrete geometric meaning to the constant α in (NSP), connecting it with the diameter of a section of the ℓ_1 ball by the null space of the matrix A (see e.g. Kashin and Temlyakov [2007] for more details).

Lemma 1.6. Given a matrix $A \in \mathbb{R}^{n \times p}$, denote

$$\frac{1}{2} \mathbf{diam}(B_1^p \cap \text{Null}(A)) = \sup_{\substack{Az=0 \\ \|z\|_1 \leq 1}} \|z\|_2,$$

the radius of the section of the ℓ_1 ball B_1^p by the null space of the matrix A and

$$s_A \triangleq 1 / \mathbf{diam}(B_1^p \cap \text{Null}(A))^2, \quad (8)$$

a recovery threshold. Then A satisfies **(NSP)** at any order $s < s_A$ with constant

$$\alpha = 2\sqrt{s_A/s} - 1 > 1. \quad (9)$$

Proof. For any $z \in \text{Null}(A)$ and support set S with $\text{Card}(S) \leq s$, using equivalence of norms and definition of the radius,

$$\|z_S\|_1 \leq \sqrt{s} \|z\|_2 \leq \frac{1}{2} \sqrt{\frac{s}{s_A}} \|z\|_1 = \frac{1}{2} \sqrt{\frac{s}{s_A}} (\|z_S\|_1 + \|z_{S^c}\|_1),$$

which means, as $s < s_A$,

$$\|z_{S^c}\|_1 \geq (2\sqrt{s_A/s} - 1) \|z_S\|_1,$$

hence the desired result. ■

With s_A defined in (8), for any signal x^* of sparsity $s < s_A$, the sharpness bound (**Sharp**) then reads

$$\|x\|_1 - \|x^*\|_1 \geq \left(1 - \sqrt{s/s_A}\right) \|x - x^*\|_1,$$

and the optimal restart scheme defined in Proposition 1.4 has complexity

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \exp\left(-\left(1 - \sqrt{s/s_A}\right) \frac{e}{2\sqrt{p}} N\right) \epsilon_0,$$

which means that, given a sensing matrix A with recovery threshold s_A , the sparser the signal, the faster the algorithm.

Precise estimates of the diameter of random sections of norm balls can be computed using classical results in geometric functional analysis. The low M^* estimates of Pajor and Tomczak-Jaegermann [1986] (see [Vershynin, 2011, Theorem 3.1] for a concise presentation) show that when $E \subset \mathbb{R}^p$ is a random subspace of codimension n (e.g. the null space of a random matrix $A \in \mathbb{R}^{n \times p}$), then

$$\mathbf{diam}(B_1^p \cap E) \leq c \sqrt{\frac{\log p}{n}},$$

with high probability, where $c > 0$ is an absolute constant. This means that the recovery threshold s_A satisfies

$$s_A \geq n / (c^2 \log p),$$

with high probability and leads to the following corollary.

Corollary 1.7. Given a random sampling matrix $A \in \mathbb{R}^{n \times p}$ and a signal x^* with sparsity $s < n / (c^2 \log p)$, (**Restart**) scheme with optimal parameters defined in (4) outputs a point \hat{y} such that

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \exp\left(-\left(1 - c \sqrt{\frac{s \log p}{n}}\right) \frac{e}{2\sqrt{p}} N\right) \epsilon_0,$$

with high probability, where c is a universal constant and N is the total number of iterations.

This means that the complexity of the optimization problem (ℓ_1 recovery) is controlled by the oversampling ratio n/s . In other words, while increasing the number of samples increases the time complexity of elementary operations of the algorithm, it also increases its rate of convergence.

2. RENEGAR'S CONDITION NUMBER AND RESTRICTED SINGULAR VALUES

We first gave concrete evidence of the link between optimization complexity and recovery performance for the exact recovery problem by highlighting sharpness properties of the objective around the true signal, given by the null space condition. We now take a step back and consider results on the underlying conic geometry of recovery problems that also control both computational and statistical aspects.

On the statistical side, minimal conically restricted singular values are known to control recovery performance in robust recovery problems. On the computational side, Renegar's condition number, a well known computational complexity measure for conic convex programs, controls the cost of obtaining optimality certificates for exact recovery and the sharpness of exact recovery problems (hence computational complexity of the (Restart) scheme presented in the previous section). Numerical experiments will then illustrate its relevance to control numerous other classical algorithms. By observing that minimal conically restricted singular values match the worst case of Renegar's condition number on sparse signals, our analysis shows once more that one single geometrical quantity controls both statistical robustness and computational complexity of recovery problems.

2.1. Conic linear systems. Conic linear systems arise naturally from optimality conditions of the exact recovery problem. To see this, define the tangent cone at point x with respect to the ℓ_1 norm, that is, the set of descent directions for $\|\cdot\|_1$ at x , as

$$\mathcal{T}(x) = \text{cone}\{z : \|x + z\|_1 \leq \|x\|_1\}. \quad (10)$$

As shown for example by Chandrasekaran et al. [2012, Prop 2.1] a point x is then the unique optimum of the exact recovery problem (ℓ_1 recovery) if and only if $\text{Null}(A) \cap \mathcal{T}(x) = \{0\}$, that is, there is no point satisfying the linear constraints that has lower ℓ_1 norm than x . Correct recovery of an original signal x^* is therefore certified by the infeasibility of a conic linear system of the form

$$\begin{aligned} & \text{find } z \\ & \text{s.t. } Az = 0 \\ & \quad z \in C, z \neq 0, \end{aligned} \quad (\mathbf{P}_{A,C})$$

where C is a closed convex cone and A a given matrix. For both computational and statistical aspects we will be interested in the distance to feasibility. On the computational side this will give a distance to ill-posedness that plays the role of a condition number. On the statistical side it will measure the amount of perturbation that the recovery can handle.

Definition 2.1 (Distance to feasibility). *Writing $\mathcal{M}_C = \{A \in \mathbb{R}^{n \times p} : (\mathbf{P}_{A,C}) \text{ is infeasible}\}$, distance to feasibility is defined as*

$$\sigma_C(A) \triangleq \inf_{\Delta A} \{\|\Delta A\|_2 : A + \Delta A \notin \mathcal{M}_C\}. \quad (11)$$

A geometric analysis of the problem explicits the distance to feasibility in terms of minimal conically restricted singular value, as recalled in the following lemma.

Lemma 2.2. *Given a matrix $A \in \mathbb{R}^{p \times n}$ and a closed convex cone C , the distance to feasibility of $(\mathbf{P}_{A,C})$ is given by*

$$\sigma_C(A) = \min_{\substack{x \in C \\ \|x\|_2=1}} \|Ax\|_2. \quad (12)$$

Proof. We recall the short proof of Amelunxen and Lotz [2014, Lemma 3.2]. Similar results have been derived by Freund and Vera [1999b, Theorem 2] and Belloni and Freund [2009, Lemma 3.2]. Let $z \in C$, with $\|z\|_2 = 1$, achieve the minimum above. Then $\Delta A = -Az z^T$ satisfies $(A + \Delta A)z = 0$, so $A + \Delta A \notin \mathcal{M}_C$ and

$$\sigma_C(A) \leq \|\Delta A\|_2 = \|Az\|_2 \|z\|_2 = \min_{\substack{x \in C \\ \|x\|_2=1}} \|Ax\|_2.$$

On the other hand denote ΔA a perturbation such that $A + \Delta A \notin \mathcal{M}_C$. Then there exists $z \in C \setminus \{0\}$ such that $(A + \Delta A)z = 0$. Thus we have

$$\|\Delta A\|_2 \geq \frac{\|\Delta Az\|_2}{\|z\|_2} = \frac{\|Az\|_2}{\|z\|_2} \geq \min_{\substack{x \in C \\ \|x\|_2=1}} \|Ax\|_2.$$

Taking the infimum on the left-hand side over all ΔA such that $A + \Delta A \notin \mathcal{M}_C$ concludes the proof. ■

Expression (12) writes distance to infeasibility as a cone restricted eigenvalue. Minimal cone restricted eigenvalues also directly characterize recovery performance as we recall now.

2.2. Recovery performance of robust recovery. The previous section showed that recovery of a signal x^* is ensured by infeasibility of the conic linear system $(P_{A, \mathcal{T}(x^*)})$, i.e. positiveness of the minimal conically restricted singular value $\sigma_{\mathcal{T}(x^*)}(A)$. We now show how this quantity also controls recovery performance in the presence of noise. In that case, the robust recovery problem attempts to retrieve an original signal x^* by solving

$$\begin{aligned} & \text{minimize} && \|x\|_1 \\ & \text{subject to} && \|Ax - b\|_2 \leq \delta \|A\|_2, \end{aligned} \quad (\text{Robust } \ell_1 \text{ recovery})$$

in the variable $x \in \mathbb{R}^p$, with the same design matrix $A \in \mathbb{R}^{n \times p}$, where $b \in \mathbb{R}^n$ are given observations perturbed by noise of level $\delta > 0$. The following classical result then bounds reconstruction error in terms of $\sigma_{\mathcal{T}(x^*)}(A)$.

Lemma 2.3. *Given a coding matrix $A \in \mathbb{R}^{n \times p}$ and an original signal x^* , suppose we observe $b = Ax^* + w$ where $\|w\|_2 \leq \delta \|A\|_2$ and denote an optimal solution of (Robust ℓ_1 recovery) by \hat{x} . If the minimal singular value $\sigma_{\mathcal{T}(x^*)}(A)$ in (12) restricted to the tangent cone $\mathcal{T}(x^*)$ in (10) is positive, the following error bound holds:*

$$\|\hat{x} - x^*\|_2 \leq 2 \frac{\delta \|A\|_2}{\sigma_{\mathcal{T}(x^*)}(A)}. \quad (13)$$

Proof. We recall the short proof of Chandrasekaran et al. [2012, Prop. 2.2]. Both \hat{x} and x^* are feasible for (Robust ℓ_1 recovery) and \hat{x} is optimal, so that $\|\hat{x}\|_1 \leq \|x^*\|_1$. Thus, the error vector $\hat{x} - x^*$ is in the tangent cone $\mathcal{T}(x^*)$. By the triangle inequality,

$$\|A(\hat{x} - x^*)\|_2 \leq \|A\hat{x} - b\|_2 + \|Ax^* - b\|_2 \leq 2\delta \|A\|_2.$$

Furthermore, by definition of $\sigma_{\mathcal{T}(x^*)}(A)$,

$$\|A(\hat{x} - x^*)\|_2 \geq \sigma_{\mathcal{T}(x^*)}(A) \|\hat{x} - x^*\|_2.$$

Combining the two concludes the proof. ■

Therefore the robustness of the coding matrix A on all s -sparse signals is measured by

$$\mu_s(A) \triangleq \inf_{x: \|x\|_0 \leq s} \min_{\substack{z \in \mathcal{T}(x) \\ \|z\|_2=1}} \|Az\|_2. \quad (14)$$

Expression of this minimal conically singular value can be simplified by identifying the tangent cones on s -sparse signals, as done in the following lemma.

Lemma 2.4. For any subset $S \subset \llbracket 1, p \rrbracket$, let

$$\mathcal{E}_S = \{z : \|z_{S^c}\|_1 \leq \|z_S\|_1\} \quad \text{and} \quad \mathcal{F}_S = \bigcup_{x: x=x_S} \mathcal{T}(x),$$

then $\mathcal{E}_S = \mathcal{F}_S$.

Proof. Let $z \in \mathcal{E}_S$, take $x = -z_S$, then

$$\|x + z\|_1 = \|z_{S^c}\|_1 \leq \|z_S\|_1 = \|x\|_1.$$

Therefore $z \in \mathcal{T}(x) \subset \mathcal{F}_S$ as $x = x_S$.

Conversely let $z \in \mathcal{F}_S$, and $x \in \mathbb{R}^p$, with $x = x_S$, such that $z \in \mathcal{T}(x)$. Then

$$\|x + z\|_1 = \|x + z_S\|_1 + \|z_{S^c}\|_1 \geq \|x\|_1 - \|z_S\|_1 + \|z_{S^c}\|_1.$$

As $z \in \mathcal{T}(x)$, this implies $\|z_{S^c}\|_1 \leq \|z_S\|_1$, so $z \in \mathcal{E}_S$ and we conclude that $\mathcal{E}_S = \mathcal{F}_S$. ■

Therefore, the previous expression for the minimal conically restricted singular value (14) can be equivalently stated as

$$\mu_s(A) = \min_{\substack{S \subset \llbracket 1, p \rrbracket \\ \text{Card}(S) \leq s}} \min_{\substack{\|z_{S^c}\|_1 \leq \|z_S\|_1 \\ \|z\|_2 = 1}} \|Az\|_2. \quad (15)$$

This quantity was introduced in [Bickel et al., 2009] and further explored in e.g. [Van De Geer and Bühlmann, 2009]. Bickel et al. [2009] notably showed that it controls estimation performance of LASSO and Dantzig selector. Observe that positiveness of $\mu_s(A)$ is equivalent to (NSP) at order s with constant 1 which shows necessity of (NSP) for sparse recovery.

Since both null space property and conically restricted singular values are necessary and sufficient conditions for exact recovery they may have been linked previously in the literature. Here we derive estimates for the constant in (NSP) from the minimal cone restricted singular value using tools from conic linear systems. We search for α such that (NSP) is satisfied at order s . Equivalently we search for α such that for any support S of cardinality at most s , the conic linear system

$$\begin{aligned} \text{find } & z \\ \text{s.t. } & Az = 0 \\ & \|z_{S^c}\|_1 \leq \alpha \|z_S\|_1, z \neq 0 \end{aligned} \quad (16)$$

is infeasible. Notice that system (16) for $\alpha > 1$ is a perturbed version of the case $\alpha = 1$, so the problem reduces to studying the sensitivity to perturbations of conic linear systems as shown in the following lemma.

Lemma 2.5. Given a matrix $A \in \mathbb{R}^{n \times p}$ and an integer $s \in \llbracket 1, p \rrbracket$, if the minimal conically restricted singular value $\mu_s(A)$ in (14) and (15) is positive, then A satisfies (NSP) at order s for any constant

$$\alpha \leq \left(1 - \frac{\mu_s(A)}{\|A\|_2}\right)^{-1}.$$

Proof. For a support S of cardinality at most s , write P the orthogonal projector on this support (that is, $Px = x_S$), $\bar{P} = \mathbf{I} - P$ its orthogonal projector and define the closed convex cone $C_S = \{z : \|z_{S^c}\|_1 \leq \|z_S\|_1\}$. Given $\alpha \geq 1$, denote $H = \alpha^{-1}P + \bar{P} = \mathbf{I} - (1 - \alpha^{-1})P$. Observe that

$$\{z : \|z_{S^c}\|_1 \leq \alpha \|z_S\|_1\} = HC_S.$$

Therefore, the conic linear system (16) reads

$$\begin{aligned} \text{find } & z \\ \text{s.t. } & Az = 0 \\ & z \in HC_S, z \neq 0. \end{aligned}$$

As H is invertible, this is equivalent to

$$\begin{aligned} & \text{find } z \\ & \text{s.t. } AH z = 0 \\ & \quad z \in C_S, z \neq 0. \end{aligned} \tag{17}$$

Therefore, if the conic linear system

$$\begin{aligned} & \text{find } z \\ & \text{s.t. } Az = 0 \\ & \quad z \in C_S, z \neq 0 \end{aligned}$$

is infeasible, that is $\sigma_{C_S}(A) > 0$, by Lemma 2.2, which is true for $\mu_s > 0$, then by definition of the distance to feasibility, (17) is also infeasible provided $\|AH - A\|_2 \leq \sigma_{C_S}(A)$, which holds for any $\alpha \geq 1$ such that

$$(1 - \alpha^{-1})\|AP\|_2 \leq \sigma_{C_S}(A).$$

Using that $\|AP\|_2 \leq \|A\|_2$, infeasibility is ensured in particular for any α such that

$$1 - \frac{\sigma_{C_S}(A)}{\|A\|_2} \leq \alpha^{-1}.$$

To ensure infeasibility of the conic linear systems (16) for any support S , it suffices to take α such that

$$1 - \frac{\mu_s(A)}{\|A\|_2} \leq \alpha^{-1}.$$

This means that (NSP) at order s is satisfied for any

$$\alpha \leq \left(1 - \frac{\mu_s(A)}{\|A\|_2}\right)^{-1}$$

where we used that, by definition of the minimal conically restricted singular value, $\mu_s(A) \leq \|A\|_2$ (in case of equality (NSP), will be satisfied for any $\alpha \geq 1$). ■

We now relate the minimal cone restricted singular value to computational complexity measures.

2.3. Computational complexity of recovery problems. Computational complexity for convex optimization problems is often described in terms of polynomial functions of problem size. This produces a clear link between problem structure and computational complexity but fails to account for the nature of the data. If we use linear systems as a basic example, unstructured linear systems of dimension n can be solved with complexity $O(n^3)$ regardless of the matrix values, but iterative solvers will converge much faster on systems that are better conditioned. The seminal work of Renegar [1995b, 2001] extends this notion of conditioning to optimization problems, producing data-driven bounds on the complexity of solving conic programs, and showing that the number of outer iterations of interior point algorithms increases as the distance to ill-posedness decreases.

2.3.1. Renegar's condition number. Renegar's condition number [Renegar, 1995b,a; Peña, 2000] provides a data-driven measure of the complexity of certifying infeasibility of a conic linear system of the form presented in (P_{A,C}) (the larger the condition number, the harder the problem). It is rooted in the sensible idea that certifying infeasibility is easier if the problem is far from being feasible. It is defined as the scale invariant reciprocal of the distance to feasibility $\sigma_C(A)$, defined in (11), of problem (P_{A,C}), i.e.

$$\mathcal{R}_C(A) \triangleq \frac{\|A\|_2}{\sigma_C(A)} = \|A\|_2 / \min_{\substack{x \in C \\ \|x\|_2=1}} \|Ax\|_2. \tag{18}$$

Notice that, if C were the whole space \mathbb{R}^p , and if $A^T A$ were full-rank (never the case if $n < p$), then $\sigma_C(A)$ would be the smallest singular value of A . As a result, $\mathcal{R}_C(A)$ would reduce to the classical condition number of A (and to ∞ when $A^T A$ is rank-deficient). Renegar's condition number is necessarily smaller

(better) than the latter, as it further incorporates the notion that A need only be well conditioned along those directions that matter with respect to C .

2.3.2. Complexity of certifying optimality. In a first step, we study the complexity of the oracle certifying optimality of a candidate solution x to (ℓ_1 recovery) as a proxy for the problem of computing an optimal solution to this problem. As mentioned in Section 2.1, optimality of a point x is equivalent to infeasibility of

$$\begin{aligned} & \text{find } z \\ & \text{s.t. } Az = 0 \\ & \quad z \in \mathcal{T}(x), z \neq 0, \end{aligned} \tag{P_{A,\mathcal{T}(x)}}$$

where the tangent cone $\mathcal{T}(x)$ is defined in (10). By a theorem of alternative, infeasibility of $(P_{A,\mathcal{T}(x)})$ is equivalent to feasibility of the dual problem

$$\begin{aligned} & \text{find } y \\ & \text{s.t. } A^T y \in \text{int}(\mathcal{T}(x)^\circ), \end{aligned} \tag{D_{A,\mathcal{T}(x)}}$$

where $\mathcal{T}(x)^\circ$ is the polar cone of $\mathcal{T}(x)$. Therefore, to certify infeasibility of $(P_{A,\mathcal{T}(x)})$ it is sufficient to exhibit a solution for the dual problem $(D_{A,\mathcal{T}(x)})$.

Several references have connected Renegar's condition number and the complexity of solving such conic linear systems using various algorithms [Renegar, 1995b; Freund and Vera, 1999a; Epelman and Freund, 2000; Renegar, 2001; Vera et al., 2007; Belloni et al., 2009]. In particular, Vera et al. [2007] linked it to the complexity of solving the primal dual pair $(P_{A,\mathcal{T}(x)})$ – $(D_{A,\mathcal{T}(x)})$ using a barrier method. They show that the number of outer barrier method iterations grows as

$$O(\sqrt{\rho} \log(\rho \mathcal{R}_{\mathcal{T}(x)}(A))),$$

where ρ is the barrier parameter, while the conditioning (hence the complexity) of the linear systems arising at each interior point iteration is controlled by $\mathcal{R}_{\mathcal{T}(x)}(A)^2$. This link was also tested empirically on linear programs using the NETLIB library of problems by Ordóñez and Freund [2003], where computing times and number of iterations were regressed against estimates of the condition number computed using the approximations for Renegar's condition number detailed by Freund and Vera [2003].

Studying the complexity of computing an optimality certificate gives insights on the performance of oracle based optimization techniques such as the ellipsoid method. We now show how Renegar's condition also controls the number steps in the (Restart) scheme presented in Section 1.2.

2.3.3. Complexity of restart scheme with Renegar's condition number. Convergence of the (Restart) scheme presented in Section 1.2 is controlled by the sharpness of the problem deduced from (NSP). We now observe that sharpness is controlled by the worst case Renegar condition number for the optimality certificates $(P_{A,\mathcal{T}(x)})$ on all s -sparse signals, defined as

$$\mathcal{R}_s(A) \triangleq \sup_{x: \|x\|_0 \leq s} \mathcal{R}_{\mathcal{T}(x)}(A) = \|A\|_2 / \mu_s(A). \tag{19}$$

Connecting Lemmas 2.3, 2.5 and Proposition 1.4 we get the following corollary.

Corollary 2.6. *Given a coding matrix $A \in \mathbb{R}^{n \times p}$ and a sparsity level $s \geq 1$, if $\mathcal{R}_s(A) < +\infty$ in (19) then optimal (Restart) scheme achieves an ϵ precision in at most*

$$O((2\mathcal{R}_s(A) - 1) \log \epsilon^{-1})$$

iterations.

This shows that Renegar's condition number explicitly controls the convergence of an algorithmic scheme devoted to the exact recovery problem (ℓ_1 recovery), through its link with sharpness.

On the statistical side, we observed that the minimal conically restricted singular value controls recovery performance of robust procedures and that its positivity ensures exact recovery. On the computational side,

we presented the role of Renegar’s condition number as a computational complexity measure for sparse recovery problems. A key observation is that the worst case of Renegar’s condition number $\mathcal{R}_s(A)$, defined in (19), matches the minimal conically restricted singular value defined in (14). Once again, a single quantity controls both aspects. This at least partially explains the common empirical observation (see, e.g., [Donoho and Tsaig \[2008\]](#)) that problem instances where statistical estimation succeeds are computationally easier to solve.

2.4. Computational complexity for inexact recovery. When the primal problem $(\mathbf{P}_{A,\mathcal{T}(x)})$ is feasible, so that $\sigma_{\mathcal{T}(x)}(A) = 0$, Renegar’s condition number as defined here is infinite. While this correctly captures the fact that, in that regime, statistical recovery does not hold, it does not properly capture the fact that, when $(\mathbf{P}_{A,\mathcal{T}(x)})$ is “comfortably” feasible, certifying so is easy, and algorithms terminate quickly (although they return a useless estimator). From both a statistical and a computational point of view, the truly delicate cases correspond to problem instances for which both $(\mathbf{P}_{A,\mathcal{T}(x)})$ and $(\mathbf{D}_{A,\mathcal{T}(x)})$ are only barely feasible or infeasible. This is illustrated in simple numerical example by [Boyd and Vandenberghe \[2004, §11.4.3\]](#) and in our numerical experiments, corresponding to the peaks in the CPU time plots of the right column in [Figure 4](#): problems where sparse recovery barely holds/fails are relatively harder. For simplicity, we only focused here on distance to feasibility for problem $(\mathbf{P}_{A,\mathcal{T}(x)})$. However, it is possible to symmetrize the condition numbers used here as described by [Amelunxen and Lotz \[2014, §1.3\]](#), where a symmetric version of the condition number is defined as

$$\bar{\mathcal{R}}_{\mathcal{T}(x)}(A) = \min \left\{ \frac{\|A\|}{\sigma_{\mathcal{T}(x)}^P(A)}, \frac{\|A\|}{\sigma_{\mathcal{T}(x)}^D(A)} \right\},$$

where $\sigma_{\mathcal{T}(x)}^P(A)$ and $\sigma_{\mathcal{T}(x)}^D(A)$ denote the distance to feasibility of respectively $(\mathbf{P}_{A,\mathcal{T}(x)})$ and $(\mathbf{D}_{A,\mathcal{T}(x)})$. This quantity peaks for programs that are nearly feasible/infeasible.

As we noticed in [Section 1.2](#), a Łojasiewicz inequality (6) for the (ℓ_1 recovery) problem is sufficient to ensure linear convergence of the restart scheme. Connecting the symmetrized Renegar condition number to the Łojasiewicz inequality constant γ may then produce complexity bounds for the restart scheme beyond the recovery case. Łojasiewicz inequalities for convex programs have indeed proven their relevance. They were used by [Fercq and Qu \[2016\]](#); [Roulet and d’Aspremont \[2017\]](#) to accelerate classical methods, in particular on the LASSO problem. Lower computational bounds for convex optimization problems satisfying it were also studied by [Nemirovskii and Nesterov \[1985, Page 6\]](#). Although the Łojasiewicz inequality is proven to be satisfied by a broad class of functions [[Bolte et al., 2007](#)], quantifying its parameters is still a challenging problem that would enable better parameter choices for appropriate algorithms.

2.5. Other algorithms. The restart scheme presented in [Section 1.2](#) is of course not the only one to solve problem (ℓ_1 recovery) in practice and it has not been analyzed in the noisy case. However, we will observe in the numerical experiments of [Section 4](#) that the condition number is correlated with the empirical performance of efficient recovery algorithms such as LARS [[Efron et al., 2004](#)] and Homotopy [[Donoho and Tsaig, 2008](#); [Asif and Romberg, 2014](#)]. On paper, the computational complexities of (ℓ_1 recovery) and (**Robust ℓ_1 recovery**) are very similar (in fact, infeasible start primal-dual algorithms designed for solving (ℓ_1 recovery) actually solve problem (**Robust ℓ_1 recovery**) with δ small). However in our experiments, we did observe sometimes significant differences in behavior between the noisy and noiseless case.

3. GENERALIZATION TO COMMON SPARSITY INDUCING NORMS

In this section we generalize previous results to sparse recovery problems in (non-overlapping) group norms or nuclear norm. Group norms arise in contexts such as genomics to enforce the selection of groups of genes (e.g., [Obozinski et al. \[2011\]](#) and references therein.) The nuclear norm is used for low-rank estimation (e.g., [Recht et al. \[2008\]](#) and references therein.) We use the framework of decomposable norms

introduced by [Negahban et al. \[2009\]](#) which applies to these norms. This allows us to generalize the null space property and to derive corresponding sharpness bounds for the exact recovery problem in a broader framework. We then again relate recovery performance and computational complexity of these recovery problems.

3.1. Decomposable norms. Sparsity inducing norms have been explored from various perspectives. Here, we use the framework of decomposable norms by [Negahban et al. \[2009\]](#) to generalize our results from ℓ_1 norms to non-overlapping group norms and nuclear norms in a concise form. We then discuss the key geometrical properties of these norms and potential characterization of their conic nature.

We first recall the definition of decomposable norms by [Negahban et al. \[2009\]](#) in terms of projectors.

Definition 3.1. Decomposable norms *Given a Euclidean space E , a norm $\|\cdot\|$ on E is said to be decomposable if there exists a family of orthogonal projectors \mathcal{P} such that*

- (i) *to each $P \in \mathcal{P}$ is associated a non-negative weight $\eta(P)$ and an orthogonal projector \bar{P} such that $P\bar{P} = \bar{P}P = 0$, and*
- (ii) *for any $x \in E$ and $P \in \mathcal{P}$, $\|Px + \bar{P}x\| = \|Px\| + \|\bar{P}x\|$.*

A signal x is then said to be s -sparse if there exists $P \in \mathcal{P}$, such that $\eta(P) \leq s$ and $Px = x$.

We now detail the family of projectors for some decomposable norms of interest.

ℓ_1 norm. In the the ℓ_1 norm case, $E = \mathbb{R}^p$ and \mathcal{P} is the set of projectors on coordinate subspaces of \mathbb{R}^p , that is, \mathcal{P} contains all projectors which zero out all coordinates of a vector except for a subset of them, which are left unaffected. The maps \bar{P} are the complementary projectors: $\bar{P} = \mathbf{I} - P$. Property (ii) is the classical decomposability of the ℓ_1 norm. Naturally, the complexity level corresponds to the number of coordinates preserved by P , i.e., $\nu(P) = \mathbf{Rank}(P)$. These definitions recover the usual notion of sparsity.

Group norms. Given a partition G of $\llbracket 1, p \rrbracket$ in (non-overlapping) groups $g \subset \llbracket 1, p \rrbracket$, the group norm is defined for $x \in \mathbb{R}^p$ as

$$\|x\| = \sum_{g \in G} \|x_g\|_r,$$

where $\|x_g\|_r$ is the ℓ_r -norm of the projection of x onto the coordinates defined by g . The cases $r = 2, \infty$ correspond respectively to ℓ_1/ℓ_2 and ℓ_1/ℓ_∞ block norms. Here, $E = \mathbb{R}^p$ and the family \mathcal{P} is composed of orthogonal projectors onto coordinates defined by (disjoint) unions of groups g , and $\bar{P} = \mathbf{I} - P$. Formally, to each P we associate $F \subset G$ such that for any $x \in E$, $(Px)_g = x_g$ if $g \in F$ and $(Px)_g = 0$ otherwise. Decomposability (ii) then clearly holds. To each group g we associate a weight η_g and for a projector $P \in \mathcal{P}$ with associated $F \subset G$, $\eta(P) = \sum_{g \in F} \eta_g$. A classical choice of weights is $\eta_g = 1$ for all $g \in G$.

Nuclear norm. The nuclear norm is defined for matrices $X \in \mathbb{R}^{p \times q}$ with singular values $\sigma_i(X)$ as

$$\|X\| = \sum_{k=1}^{\min(p,q)} \sigma_k(X).$$

Here $E = \mathbb{R}^{p \times q}$ and its associated family of projectors contains P such that

$$P: X \mapsto P_{\text{left}} X P_{\text{right}},$$

and

$$\bar{P}: X \mapsto (\mathbf{I} - P_{\text{left}}) X (\mathbf{I} - P_{\text{right}}),$$

where $P_{\text{left}} \in \mathbb{R}^{p \times p}$ and $P_{\text{right}} \in \mathbb{R}^{q \times q}$ are orthogonal projectors. Their weights are defined as $\eta(P) = \max(\mathbf{Rank}(P_{\text{left}}), \mathbf{Rank}(P_{\text{right}}))$ defining therefore s -sparse matrices as matrices of rank at most s . As P and \bar{P} project on orthogonal row and column spaces, condition (ii) holds.

Decomposable norms offer a unified nomenclature for the study of sparsity inducing norms. However, they appear to be essentially restricted to the three examples presented above. Moreover, it is not clear if

their definition is sufficient to characterize the conic nature of these norms, in particular in the nuclear norm case that will require additional linear algebra results. In comparison, the framework proposed by Juditsky et al. [2014] can encompass *non-latent* overlapping groups. For future use, we simplify the third property of their definition [Juditsky et al., 2014, Section 2.1] in Appendix B. It is not clear how this view can be used for latent overlapping group norms presented by Obozinski et al. [2011] applied in biology. Moreover the sufficient conditions that Juditsky et al. [2014] present are sufficient but not necessary in the nuclear norm case. Better characterizing the key geometrical properties of these norms is therefore a challenging research direction.

3.2. Sharpness and generalized null space property. From now on, we assume that we are given an ambient Euclidean space E with one of the three decomposable norms $\|\cdot\|$ presented in previous section, i.e. ℓ_1 , group or nuclear norm, and the associated family of orthogonal projectors \mathcal{P} as introduced in Definition 3.1. We study the sparse recovery problem

$$\begin{aligned} & \text{minimize} && \|x\| \\ & \text{subject to} && A(x) = b \end{aligned} \quad (\text{Sparse recovery})$$

in the variable $x \in E$, where A is a linear operator onto \mathbb{R}^n and the observations $b \in \mathbb{R}^n$ are taken from an original point x^* such that $b = A(x^*)$. We begin by generalizing the null space property in this setting.

Definition 3.2. (Generalized Null space Property) *A linear operator A on E satisfies the Generalized Null Space Property (GNSP) for orthogonal projector $P \in \mathcal{P}$ with constant $\alpha \geq 1$ if and only if for any $z \in \text{Null}(A) \setminus \{0\}$ such that $z = Pz + \bar{P}z$,*

$$\alpha \|Pz\| < \|\bar{P}z\|. \quad (\text{GNSP})$$

The linear operator A satisfies the Generalized Null Space Property at order s with constant $\alpha \geq 1$ if it satisfies it for any P such that $\eta(P) \leq s$.

Notice that if $\bar{P} = \mathbf{I} - P$, the condition $z = Pz + \bar{P}z$ is not restrictive. However it will be useful to prove necessity of (GNSP) for the nuclear norm. In that case, observe that it is equivalent to the condition introduced by Oymak and Hassibi [2010], i.e.

$$\forall z \in \text{Null}(A) \setminus \{0\}, \quad \alpha \sum_{i=1}^s \sigma_i(z) < \sum_{i=s+1}^{\min(p,q)} \sigma_i(z),$$

where $\sigma_i(z)$ are the singular values of z in decreasing order. Notice also that we recover the classical Definition NSP in the ℓ_1 case. The sharpness bound then easily follows if $\bar{P} = \mathbf{I} - P$. In the case of the nuclear norm it requires additional linear algebra results.

Proposition 3.3. *Given a linear operator A that satisfies (GNSP) at order s with constant α , if the original point x^* is s -sparse, then for any $x \in E$ satisfying $A(x) = b$, $x \neq x^*$, we have*

$$\|x\| - \|x^*\| > \frac{\alpha - 1}{\alpha + 1} \|x - x^*\|.$$

This implies recovery, i.e., optimality of x^ for (Sparse recovery).*

Proof. Denote P such that $\eta(P) \leq s$ and $Px^* = x^*$, which defines its sparsity. Let $x \neq x^*$ such that $A(x) = b$, so $z = x - x^* \in \text{Null}(A)$ and $z \neq 0$. If $\bar{P} = \mathbf{I} - P$, $x = Px + \bar{P}x$ and using the decomposability (ii), we have

$$\begin{aligned} \|x\| &= \|Px^* + Pz\| + \|\bar{P}z\| \\ &\geq \|x^*\| - \|Pz\| + \|\bar{P}z\| \\ &= \|x^*\| + \|z\| - 2\|Pz\|. \end{aligned}$$

By using (GNSP), $\|z\| = \|Pz\| + \|\bar{P}z\| > (1 + \alpha)\|Pz\|$. The result follows by arranging the terms.

If $\|\cdot\|$ is the nuclear norm and $\bar{P} \neq \mathbf{I} - P$, as in [Oymak and Hassibi, 2010, Lemma 6], we use that (see Horn and Johnson [1990, Theorem 7.4.9.1])

$$\|x^* + z\| \geq \sum_{i=1}^{\min(p,q)} |\sigma_i(x^*) - \sigma_i(z)|,$$

where $\sigma_i(x^*), \sigma_i(z)$ denote the singular values in decreasing order of respectively x^* and z . Then, using that x^* has rank at most s ,

$$\begin{aligned} \|x\| &\geq \sum_{i=1}^s |\sigma_i(x^*) - \sigma_i(z)| + \sum_{i=s+1}^{\min(p,q)} \sigma_i(z) \\ &\geq \sum_{i=1}^s \sigma_i(x^*) - \sum_{i=1}^s \sigma_i(z) + \sum_{i=s+1}^{\min(p,q)} \sigma_i(z) \\ &= \|x^*\| - \|Qz\| + \|\bar{Q}z\|, \end{aligned}$$

where Q is the projector on the s largest singular directions of z and therefore \bar{Q} the projector on the $n - s$ others. These can be defined using the singular value decomposition of z such that $z = Qz + \bar{Q}z$. Then, using (GNSP) and the decomposability (ii) concludes the proof as above. ■

This shows that the sharpness bound of the form (Sharp) generalizes to non-overlapping group norms and the nuclear norm. Proposition 1.3 can also be generalized directly to this case with our definition of (GNSP). The smoothing argument and restart schemes developed in Section 1.2 can then be applied with similar linear convergence rates that essentially depend on the sharpness constant. By looking at the diameter of the section of the unit ball of the norm by the null space of A , one may also show that the oversampling ratio controls the sharpness bound as in Section 1.3.

As in Section 2, we now study the conic quantities which control statistical and optimization aspects.

3.3. Robust recovery performance and computational complexity. In this section, for a Euclidean space E and $x \in E$ we denote $\|x\|_2$ the ℓ_2 norm of its coefficients, if E is a matrix space $\|x\|_2$ is then the Frobenius norm of x .

3.3.1. Generalized cone restricted singular value. We begin by addressing the recovery performance of robust sparse recovery problems that reads

$$\begin{aligned} &\text{minimize} && \|x\| \\ &\text{subject to} && \|A(x) - b\|_2 \leq \delta \|A\|_2, \end{aligned} \quad (\text{Robust sparse recovery})$$

in the variable $x \in E$, with the same linear operator A , where the observations $b \in \mathbb{R}^n$ are affected by noise of level $\delta > 0$. For a linear operator A from E to \mathbb{R}^n , we denote its operator norm with respect to $\|\cdot\|$, $\|A\|_2 = \sup_{x \in E: \|x\|_2 \leq 1} \|A(x)\|_2$.

The results of Section 2.2 transpose directly to the general case by replacing $\|\cdot\|_1$ by $\|\cdot\|$. Precisely, assuming that $b = Ax^* + w$ where $\|w\|_2 \leq \delta \|A\|_2$, an optimal solution \hat{x} of problem (Robust sparse recovery) satisfies the error bound

$$\|\hat{x} - x^*\|_2 \leq 2 \frac{\delta \|A\|_2}{\sigma_{\mathcal{T}(x^*)}(A)},$$

where the tangent cone is defined as

$$\mathcal{T}(x) = \text{cone}\{z : \|x + z\| \leq \|x\|\},$$

and robust recovery of s -sparse signals is therefore controlled by

$$\mu_s(A) = \inf_{P \in \mathcal{P}: \eta(P) \leq s} \inf_{x \in E: Px=x} \min_{\substack{z \in \mathcal{T}(x) \\ \|z\|_2=1}} \|Az\|_2. \quad (20)$$

The key point is then to characterize the tangent cones of s -sparse signals. First, this will allow statistical estimations of $\mu_s(A)$. Second, it will enable us to estimate the constant (**GNSP**), hence sharpness of the exact recovery problem and computational complexity of associated restart schemes. This is the aim of the following lemma.

Lemma 3.4. *For a given sparsity s , write*

$$\mathcal{E} = \bigcup_{P \in \mathcal{P} : \eta(P) \leq s} \{z \in E : z = Pz + \bar{P}z, \|\bar{P}z\| \leq \|Pz\|\}$$

and

$$\mathcal{F} = \bigcup_{P \in \mathcal{P} : \eta(P) \leq s} \bigcup_{x \in E : x = Px} \mathcal{T}(x).$$

Then $\mathcal{E} = \mathcal{F}$.

Proof. Let $z \in \mathcal{E}$ and $P \in \mathcal{P}$ such that $z = Pz + \bar{P}z$. Taking $x = -Pz$ we get

$$\|x + z\| = \|\bar{P}z\| \leq \|Pz\| = \|x\|.$$

Therefore $z \in \mathcal{T}(x) \subset \mathcal{F}$. Conversely, if $z \in \mathcal{F}$, denote $x \in E$ and $P \in \mathcal{P}$ such that $x = Px$, $z \in \mathcal{T}(x)$ and $\eta(P) \leq s$. If $\bar{P} = \mathbf{I} - P$, by decomposability (ii),

$$\|x + z\| = \|Px + Pz\| + \|\bar{P}z\| \geq \|x\| - \|Pz\| + \|\bar{P}z\|.$$

Since $z \in \mathcal{T}(x)$, we have $\|x + z\| \leq \|x\|$; combined with the previous statement, this implies that $z \in \{z \in E : z = Pz + \bar{P}z, \|\bar{P}z\| \leq \|Pz\|\} \subset \mathcal{E}$. Now, if $\|\cdot\|$ is the nuclear norm, as in the proof of Proposition (3.3), we have

$$\|x + z\| \geq \|x\| - \|Qz\| + \|\bar{Q}z\|,$$

where Q is the projector on the s largest singular directions of z given by the singular value decomposition of z , so that $z = Qz + \bar{Q}z$. Therefore, $z \in \mathcal{T}(x)$ implies $z \in \{z \in E : z = Qz + \bar{Q}z, \|\bar{Q}z\| \leq \|Qz\|\} \subset \mathcal{E}$. In all cases we have therefore proven $\mathcal{E} = \mathcal{F}$. ■

Using the previous lemma, the minimal cone restricted singular value reads:

$$\mu_s(A) = \inf_{P \in \mathcal{P}, \eta(P) \leq s} \min_{\substack{z \in E, \|z\|_2 = 1 \\ z = Pz + \bar{P}z, \|\bar{P}z\| \leq \|Pz\|}} \|Az\|_2. \quad (21)$$

This quantity can then be linked to the (**GNSP**) constant, as shown in the following lemma.

Lemma 3.5. *Given a linear operator A on E , If the minimal cone restricted singular value $\mu_s(A)$, defined in (20) and reformulated in (21), is positive, then A satisfies (**GNSP**) at order s for any constant*

$$\alpha \leq \left(1 - \frac{\mu_s(A)}{\|A\|_2}\right)^{-1}.$$

Proof. For a given $P \in \mathcal{P}$, denote $C_P = \{z \in \text{Im}(P) + \text{Im}(\bar{P}) : \|\bar{P}z\| \leq \|Pz\|\}$ and define for $\alpha \geq 1$ the conic linear system

$$\begin{aligned} \text{find } & z \in \text{Im}(P) + \text{Im}(\bar{P}) \\ \text{s.t. } & A(z) = 0 \\ & \|\bar{P}z\| \leq \alpha \|Pz\|, z \neq 0. \end{aligned} \quad (22)$$

Infeasibility of this system for all $P \in \mathcal{P}$ such that $\eta(P) \leq s$ is then equivalent to (**GNSP**) at order s with constant α . Denote $H = \mathbf{I} - (1 - \alpha^{-1})P$ such that

$$\{z \in \text{Im}(P) + \text{Im}(\bar{P}) : \|\bar{P}z\| \leq \alpha \|Pz\|\} = HC_P.$$

Since H is invertible, we observe as in Lemma 2.5 that the conic linear system (22) is equivalent to

$$\begin{aligned} \text{find } & z \in \text{Im}(P) + \text{Im}(\bar{P}) \\ \text{s.t. } & A - (1 - \alpha^{-1})APz = 0 \\ & z \in C_P, z \neq 0. \end{aligned} \quad (23)$$

If this problem is infeasible for $\alpha = 1$, i.e., its distance to feasibility $\mu_{C_P}(A)$ defined in (11) is positive, then (23) is infeasible for any $\alpha \geq 1$ such that

$$(1 - \alpha^{-1})\|AP\| \leq \mu_{C_P}(A).$$

Now, if $\mu_s(A) > 0$ the conic linear system (22) will still be infeasible for any

$$\alpha \leq \left(1 - \frac{\mu_s(A)}{\|A\|_2}\right)^{-1}.$$

Thus, A satisfies (GNSP) at order s with α as above. ■

3.3.2. *Renegar's condition number.* On the computational side, denote $\mathcal{R}_{\mathcal{T}(x)}(A)$ the Renegar condition number of the conic linear system

$$\begin{aligned} \text{find } & z \\ \text{s.t. } & A(z) = 0 \\ & z \in \mathcal{T}(x), z \neq 0, \end{aligned}$$

and the worst-case Renegar condition number on s -sparse signals

$$\mathcal{R}_s(A) \triangleq \sup_{P \in \mathcal{P} : \eta(P) \leq s} \sup_{x \in E : Px = x} \mathcal{R}_{\mathcal{T}(x)}(A) = \|A\|_2 / \mu_s(A).$$

First, Renegar's condition number plays the same role as before in computing optimality certificates for the exact recovery problems. Then, combining Lemma 3.5 and Proposition 3.3 shows that the sharpness bound for exact recovery reads

$$\|x\| - \|x^*\| > \frac{1}{2\mathcal{R}_s(A) - 1} \|x - x^*\|.$$

This sharpness will then control linearly convergent restart schemes for the exact recovery problem.

Overall then, as established earlier in this paper, a single geometric quantity—namely, the minimal cone restricted singular value—appears to control both computational and statistical aspects. We now illustrate this statement on numerical experiments.

4. NUMERICAL RESULTS

In this section, we first test the empirical performance of restart schemes and its link with recovery performance. We then perform similar experiments on Renegar's condition number.

4.1. **Sharpness & restart for exact recovery.** We test the (Restart) scheme on ℓ_1 -recovery problems with random design matrices. Throughout the experiments, we use the NESTA code described in [Becker, Bobin and Candès, 2011] as the subroutine in the restart strategy. We generate a random design matrix $A \in \mathbb{R}^{n \times p}$ with i.i.d. Gaussian coefficients. We then normalize A so that $AA^T = \mathbf{I}$ (to fit NESTA's format) and generate observations $b = Ax^*$ where $x^* \in \mathbb{R}^p$ is an s -sparse vector whose nonzero coefficients are all ones. We denote \hat{x} the solution given by a common solver run at machine precision and plot convergence $f(x_t) - f^* = \|x_t\|_1 - \|\hat{x}\|_1$ (scaled such that $f(x_0) - f(\hat{x}) = 1$).

4.1.1. *Restart scheme performance.* First we compare in Figure 1 the practical scheme presented in Section 1.2.2 with a plain implementation of NESTA without restart or continuation steps. Dimensions of the problem are $p = 500$, $n = 200$ and $s = 30$. Starting from $x_0 = A^T b$, we use $\epsilon_0 = \|x_0\|_1$ as a first initial guess on the gap and perform a grid search of step size $h = 4$ for a budget of $N = 500$ iterations. The first and last schemes of the grid search were not run as they are unlikely to produce a nearly optimal restart scheme. The grid search can be parallelized and the best scheme found is plotted with a solid red line. The dashed red line represents the convergence rate accounting for the cost of the grid search. For the plain implementation of NESTA, we used different target precisions. These control indeed the smoothness of the surrogate function f_ϵ which itself controls the step size of Nesterov’s algorithm. Therefore a high precision slows down the algorithm. However for low precision NESTA can be faster but will not approximate well the original signal. Also, the theoretical bound (2) might be very pessimistic, as the surrogate function f_ϵ may approximate the ℓ_1 norm for the points of interest at a much better accuracy than ϵ .

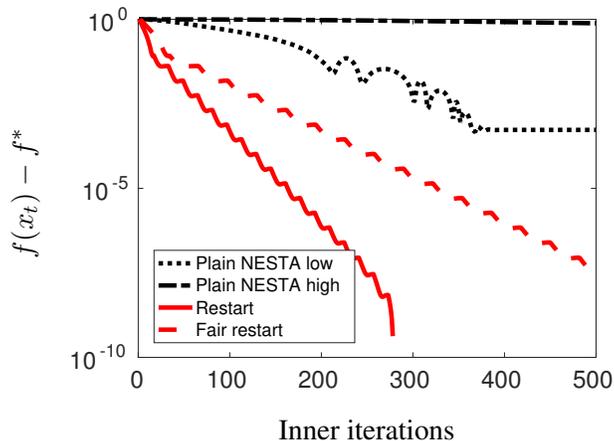


FIGURE 1. Best restarted NESTA (solid red line) and overall cost of the practical restart schemes (dashed red line) versus plain NESTA implementation with low accuracy $\epsilon = 10^{-1}$ (dotted black line) and higher accuracy $\epsilon = 10^{-3}$ (dash-dotted black line) for a budget of 500 iterations.

Overall, we observe a clear linear convergence of the restart scheme that outperforms the plain implementation. This was already observed by [Becker, Bobin and Candès \[2011\]](#) who developed their continuation steps against which we compare in Figure 2. We used default options for NESTA, namely 5 continuation steps with a stopping criterion based on the relative objective change in the surrogate function (specifically, the algorithm stops when these changes are lower than the target accuracy, set to 10^{-6}). We compare continuation steps and best restart found by grid search for different dimensions of the problem, we fix $p = 500$, $s = 30$ and vary the number of samples $n = \{200, 300\}$. Continuation steps converge faster with better conditioned problems, i.e., more samples. Otherwise they may get stuck due to the termination criterion. Notice that a lot of parameters are involved for both algorithms, in particular the target precision may play an important role, so that more extensive experiments may be needed to refine these statements.

Our goal here is to provide a simple but strong baseline with theoretical guarantees for recovery. Improving on it, as [Feroq and Qu \[2016\]](#) did for LASSO, is an appealing research direction. Sharpness may be used for example to refine the heuristic strategy of the continuations steps.

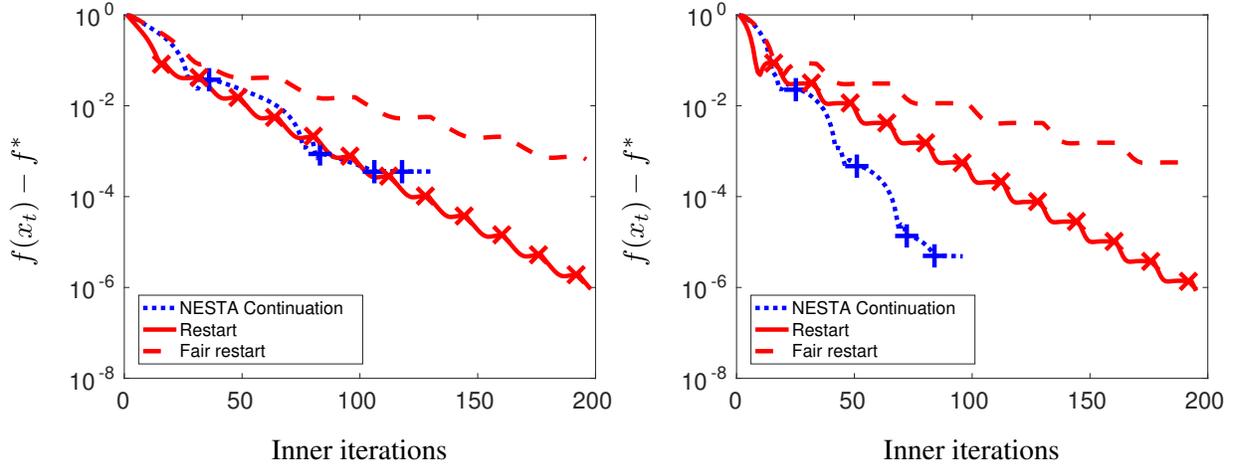


FIGURE 2. Best restarted NESTA (solid red line) and overall cost of the practical restart schemes (dashed red line) versus NESTA with 5 continuation steps (dotted blue line) for a budget of 500 iterations. Crosses represent the restart occurrences. Left: $n = 200$. Right: $n = 300$.

4.1.2. *Convergence rate and oversampling ratio.* We now illustrate the theoretical results of Section 1.3 by running the practical scheme presented in Section 1.2.2 for increasing values of the oversampling ratio $\tau = n/s$. In Figure 3, we plot the best scheme found by the grid search, that approximates the optimal scheme, for a budget of $N = 500$ iterations. We use a fine grid of step size $h = 2$. Other algorithmic parameters remain unchanged: $x_0 = A^T b$ and $\epsilon_0 = \|x_0\|_1$. We fix the dimension $p = 1000$ and either make n vary for a fixed sparsity $s = 20$ or make s vary for a fixed number of samples $n = 200$. In both cases we do observe an improved convergence for increasing oversampling ratio τ . Notice that linear convergence is observed even for $\tau < \log p$, that is, outside the recovery region. This suggests that sharpness may also hold in this case in the form of a Łojasiewicz inequality as mentioned in Sections 1.2 and 2.3.

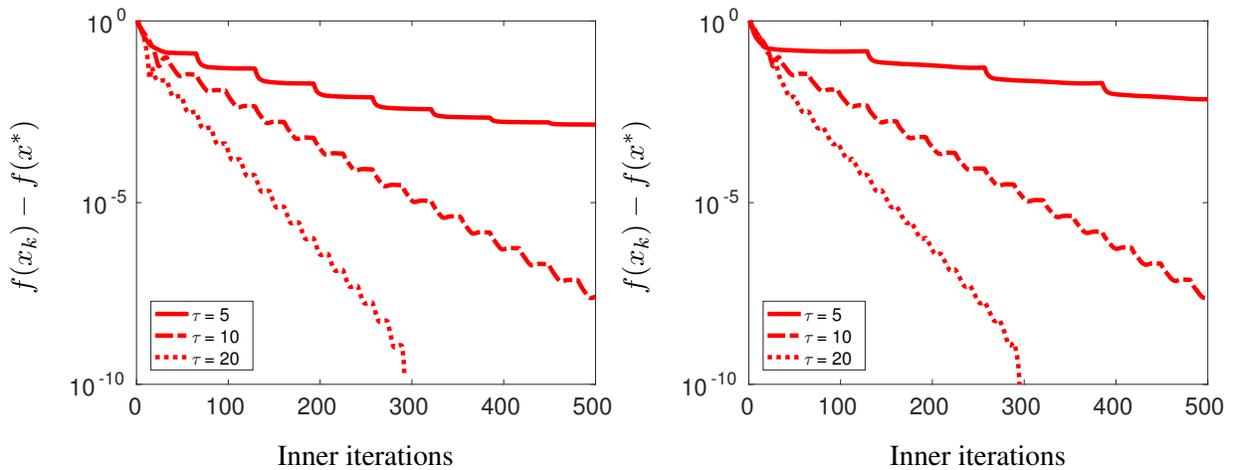


FIGURE 3. Best restart scheme found by grid search for increasing values of the oversampling ratio $\tau = n/s$. Left: sparsity $s = 20$ fixed. Right: number of samples $n = 200$ fixed.

TABLE 1. Time to achieve $\epsilon = 10^{-2}$ by the best restart scheme for increasing number of samples n

Number of samples n	100	200	400
Time in seconds for $f(x_t) - f^* < 10^{-2}$	$5.07 \cdot 10^{-2}$	$3.07 \cdot 10^{-2}$	$1.66 \cdot 10^{-2}$

In table 1, we report the time needed to achieve an $\epsilon = 10^{-2}$ precision for a fixed sparsity $s = 20$ and increasing number of samples $n = \{100, 200, 400\}$. While the cost of core operations (namely, matrix-vector multiplications) increases with n , the overall time required to reach convergence decreases. In other words, getting more samples increases the recovery performance at little or no computational cost.

4.2. Renegar’s condition number and compressed sensing performance. Our theoretical results showed that Renegar’s condition number measures the complexity for the exact recovery problem (1). However it does not a priori control convergence of the robust recovery problems defined in the introduction. This numerical section aims therefore at analyzing the relevance of this condition number for general recovery problems in the ℓ_1 case, assuming that their complexity corresponds roughly to that of checking optimality of a given point at each iteration, as mentioned in Section 2.3. We first describe how we approximate the value of $\mathcal{R}_{\mathcal{T}(x^*)}(A)$ as defined in (18) for a given original signal x^* and matrix $A \in \mathbb{R}^{n \times p}$. We then detail numerical experiments on synthetic data sets.

4.2.1. Computing $\mathcal{R}_{\mathcal{T}(x^*)}(A)$. The condition number $\mathcal{R}_{\mathcal{T}(x^*)}(A)$ appears here in upper bounds on computational complexities and statistical performances. In order to test numerically whether this quantity truly explains those features (as opposed to merely appearing in a wildly pessimistic bound), we explicitly compute it in numerical experiments.

To compute $\mathcal{R}_{\mathcal{T}(x^*)}(A)$, we propose a heuristic which computes $\sigma_{\mathcal{T}(x^*)}(A)$ in (11) and (12), the value of a nonconvex minimization problem over the cone of descent directions $\mathcal{T}(x^*)$. The closure of the latter is the polar of the cone generated by the subdifferential to the ℓ_1 -norm ball at x^* [Chandrasekaran et al., 2012, §2.3]. Let $S \subset \llbracket 1, p \rrbracket$ denote the support of x^* and $s = \text{Card}(S)$. Then, with $u = \text{sign}(x^*)$,

$$\mathcal{T}(x^*) = \text{cone} \{ z \in \mathbb{R}^p : z_S = u_S, z_{S^c} \in [-1, 1]^{p-s} \}^\circ = \{ z \in \mathbb{R}^p : \|z_{S^c}\|_1 \leq -u_S^T z_S = -u^T z \}.$$

Thus, $\sigma_{\mathcal{T}(x^*)}(A)$ is the square root of

$$\min_{z \in \mathbb{R}^p} z^T A^T A z \quad \text{s.t.} \quad \|z\|_2 = 1 \quad \text{and} \quad \|z_{S^c}\|_1 \leq -u^T z. \quad (24)$$

Let λ denote the largest eigenvalue of $A^T A$. If it were not for the cone constraint, solutions of this problem would be the dominant eigenvectors of $\lambda \mathbf{I} - A^T A$, which suggests a *projected power method* [Deshpande et al., 2014] as follows. Given an initial guess $z_0 \in \mathbb{R}^p$, $\|z_0\|_2 = 1$, iterate

$$\hat{z}_{k+1} = \text{Proj}_{\mathcal{T}(x^*)} \left((\lambda \mathbf{I} - A^T A) z_k \right), \quad z_{k+1} = \hat{z}_{k+1} / \|\hat{z}_{k+1}\|_2, \quad (25)$$

where we used the orthogonal projector to $\mathcal{T}(x^*)$,

$$\text{Proj}_{\mathcal{T}(x^*)}(\tilde{z}) = \arg \min_{z \in \mathbb{R}^p} \|z - \tilde{z}\|_2^2 \quad \text{s.t.} \quad \|z_{S^c}\|_1 \leq -u^T z. \quad (26)$$

This convex, linearly constrained quadratic program is easily solved with CVX [Grant et al., 2001]. As can be seen from KKT conditions, this iteration is a generalized power iteration [Luss and Teboulle, 2013; Journée et al., 2008]

$$z_{k+1} \in \arg \max_{z \in \mathbb{R}^p} z^T (\lambda \mathbf{I} - A^T A) z_k \quad \text{s.t.} \quad \|z\|_2 \leq 1 \quad \text{and} \quad \|z_{S^c}\|_1 \leq -u^T z.$$

From the latter, it follows that $\|Az_k\|_2$ decreases monotonically with k . Indeed, owing to convexity of $f(z) = \frac{1}{2} z^T (\lambda \mathbf{I} - A^T A) z$, we have $f(z) - f(z_k) \geq (z - z_k)^T (\lambda \mathbf{I} - A^T A) z_k$. The next iterate $z =$

z_{k+1} maximizes this lower bound on the improvement. Since $z = z_k$ is admissible, the improvement is nonnegative and $f(z_k)$ increases monotonically.

Thus, the sequence $\|Az_k\|_2$ converges, but it may do so slowly, and the value it converges to may depend on the initial iterate z_0 . On both accounts, it helps greatly to choose z_0 well. To obtain one, we modify (24) by smoothly penalizing the inequality constraint in the cost function, which results in a smooth optimization problem on the ℓ_2 sphere. Specifically, for small $\varepsilon_1, \varepsilon_2 > 0$, we use smooth proxies $h(x) = \sqrt{x^2 + \varepsilon_1^2} - \varepsilon_1 \approx |x|$ and $q(x) = \varepsilon_2 \log(1 + \exp(x/\varepsilon_2)) \approx \max(0, x)$. Then, with $\gamma > 0$ as Lagrange multiplier, we consider

$$\min_{\|z\|_2=1} \|Az\|_2^2 + \gamma \cdot q\left(u^T z + \sum_{i \in S^c} h(z_i)\right).$$

We solve the latter locally with Manopt [Boumal et al., 2014], itself with a uniformly random initial guess on the sphere, to obtain z_0 . Then, we iterate the projected power method. The value $\|Az\|_2$ is an upper bound on $\sigma_{\mathcal{T}(x^*)}(A)$, so that we obtain a lower bound on $\mathcal{R}_{\mathcal{T}(x^*)}(A)$. Empirically, this procedure, which is random only through the initial guess on the sphere, consistently returns the same value, up to five digits of accuracy, which suggests the proposed heuristic computes a good approximation of the condition number. Similarly positive results have been reported on other cones by Deshpande et al. [2014], where the special structure of the cone even made it possible to certify that this procedure indeed attains a global optimum in proposed experiments. Similarly, a generalized power method was recently shown to converge to global optimizers for the phase synchronization problem (in a certain noise regime) [Boumal, 2016; Zhong and Boumal, 2017]. This gives us confidence in the estimates produced here.

4.2.2. Sparse recovery performance. We conduct numerical experiments in the ℓ_1 case to illustrate the connection between the condition number $\mathcal{R}_{\mathcal{T}(x^*)}(A)$, the computational complexity of solving (ℓ_1 recovery), and the statistical efficiency of the estimator (**Robust ℓ_1 recovery**). Importantly, throughout the experiments, the classical condition number of A will remain essentially constant, so that the main variations cannot be attributed to the latter.

We follow a standard setup, similar to some of the experiments by Donoho and Tsaig [2008]. Fixing the ambient dimension $p = 300$ and sparsity $s = \|x^*\|_0 = 15$, we let the number of linear measurements n vary from 1 to 150. For each value of n , we generate a random signal $x^* \in \mathbb{R}^p$ (uniformly random support, i.i.d. Gaussian entries, unit ℓ_2 -norm) and a random sensing matrix $A \in \mathbb{R}^{n \times p}$ with i.i.d. standard Gaussian entries. Furthermore, for a fixed value $\delta = 10^{-2}$, we generate a random noise vector $w \in \mathbb{R}^n$ with i.i.d. standard Gaussian entries, normalized such that $\|w\|_2 = \delta \|A\|_2$, and we let $b = Ax^* + w$. This is repeated 100 times for each value of n .

For each triplet (A, x^*, b) , we first solve the noisy problem (**Robust ℓ_1 recovery**) with the L1-Homotopy algorithm ($\tau = 10^{-7}$) [Asif and Romberg, 2014], and report the estimation error $\|\hat{x} - x^*\|_2$. Then, we solve the noiseless problem (1) with L1-Homotopy and the TFOCS routine for basis pursuit ($\mu = 1$) [Becker, Candès and Grant, 2011]. Exact recovery is declared when the error is less than 10^{-5} , and we report the empirical probability of exact recovery, together with the number of iterations required by each of the solvers. The number of iterations of LARS [Efron et al., 2004] is also reported, for comparison. For L1-Homotopy, we report the computation time, normalized by the computation time required for one least-squares solve in A , as in [Donoho and Tsaig, 2008, Fig. 3], which accounts for the growth in n . Finally, we compute the classical condition number of A , $\kappa(A)$, as well as (a lower bound on) the cone-restricted condition number $\mathcal{R}_{\mathcal{T}(x^*)}(A)$, as per the previous section. As it is the computational bottleneck of the experiment, it is only computed for 20 of the 100 repetitions.

The results of Figure 4 show that the cone-restricted condition number explains both the computational complexity of (ℓ_1 recovery) and the statistical complexity of (**Robust ℓ_1 recovery**): fewer samples mean bad conditioning which in turn implies high computational complexity. We caution that our estimate of $\mathcal{R}_{\mathcal{T}(x^*)}(A)$ is only a lower bound. Indeed, for small n , the third plot on the left shows that, even in the absence of noise, recovery of x^* is not achieved by (**Robust ℓ_1 recovery**). Lemma 2.3 then requires $\mathcal{R}_{\mathcal{T}(x^*)}(A)$

to be infinite. But the computational complexity of solving (ℓ_1 recovery) is visibly favorable for small n , where far from the phase transition, problem $(\mathbf{P}_{A, \mathcal{T}(x)})$ is far from infeasibility, which is just as easy to verify as it is to certify that $(\mathbf{P}_{A, \mathcal{T}(x)})$ is infeasible when n is comfortably larger than needed. This phenomenon is best explained using a symmetric version of the condition number [Amelunxen and Lotz, 2014] (omitted here to simplify computations).

We also solved problem (ℓ_1 recovery) with interior point methods (IPM) via CVX. The number of iterations appeared mostly constant throughout the experiments, suggesting that the practical implementation of such solvers renders their complexity mostly data agnostic in the present setting. Likewise, the computation time required by L1-Homotopy on the noisy problem (Robust ℓ_1 recovery), normalized by the time of a least-squares solve, is mostly constant (at about 150). This hints that the link between computational complexity of (ℓ_1 recovery) and (Robust ℓ_1 recovery) remains to be fully explained.

Acknowledgements. VR and AA are at the Département d’Informatique at École Normale Supérieure, 2 rue Simone Iff, 75012 Paris, France. INRIA, Sierra project-team, PSL Research University. The authors would like to acknowledge support from a starting grant from the European Research Council (ERC project SIPA), an AMX fellowship, an NSF grant DMS-1719558, as well as support from the chaire *Économie des nouvelles données*, the *data science* joint research initiative with the *fonds AXA pour la recherche*.

REFERENCES

- Agarwal, A., Negahban, S. and Wainwright, M. [2011], ‘Fast global convergence of gradient methods for high-dimensional statistical recovery’, *Arxiv preprint arXiv:1104.4824*.
- Amelunxen, D. and Lotz, M. [2014], ‘Gordon’s inequality and condition numbers in conic optimization’, *arXiv preprint arXiv:1408.3016*.
- Amelunxen, D., Lotz, M., McCoy, M. B. and Tropp, J. A. [2014], ‘Living on the edge: Phase transitions in convex programs with random data’, *Information and Inference* p. iau005.
- Asif, M. S. and Romberg, J. [2014], ‘Sparse recovery of streaming signals using l1-homotopy’, *Signal Processing, IEEE Transactions on* **62**(16), 4209–4223.
- Beck, A. and Teboulle, M. [2009], ‘A fast iterative shrinkage-thresholding algorithm for linear inverse problems’, *SIAM Journal on Imaging Sciences* **2**(1), 183–202.
- Becker, S., Bobin, J. and Candès, E. J. [2011], ‘Nesta: a fast and accurate first-order method for sparse recovery’, *SIAM Journal on Imaging Sciences* **4**(1), 1–39.
- Becker, S. R., Candès, E. J. and Grant, M. C. [2011], ‘Templates for convex cone problems with applications to sparse signal recovery’, *Mathematical Programming Computation* **3**(3), 165–218.
- Belloni, A. and Freund, R. M. [2009], ‘A geometric analysis of reegar’s condition number, and its interplay with conic curvature’, *Mathematical programming* **119**(1), 95–107.
- Belloni, A., Freund, R. M. and Vempala, S. [2009], ‘An efficient rescaled perceptron algorithm for conic systems’, *Mathematics of Operations Research* **34**(3), 621–641.
- Bickel, P., Ritov, Y. and Tsybakov, A. [2009], ‘Simultaneous analysis of lasso and dantzig selector’, *The Annals of Statistics* **37**(4), 1705–1732.
- Bolte, J., Daniilidis, A. and Lewis, A. [2007], ‘The lojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems’, *SIAM Journal on Optimization* **17**(4), 1205–1223.
- Boumal, N. [2016], ‘Nonconvex phase synchronization’, *SIAM Journal on Optimization* **26**(4), 2355–2377.
- Boumal, N., Mishra, B., Absil, P.-A. and Sepulchre, R. [2014], ‘Manopt, a Matlab toolbox for optimization on manifolds’, *Journal of Machine Learning Research* **15**, 1455–1459.
URL: <http://www.manopt.org>
- Boyd, S. and Vandenberghe, L. [2004], *Convex Optimization*, Cambridge University Press.
- Candès, E. and Tao, T. [2006], ‘Near-optimal signal recovery from random projections: Universal encoding strategies?’, *IEEE Transactions on Information Theory* **52**(12), 5406–5425.
- Chandrasekaran, V. and Jordan, M. I. [2013], ‘Computational and statistical tradeoffs via convex relaxation’, *Proceedings of the National Academy of Sciences* **110**(13), 1181–1190.

- Chandrasekaran, V., Recht, B., Parrilo, P. and Willsky, A. [2012], ‘The convex geometry of linear inverse problems’, *Foundations of Computational Mathematics* **12**(6), 805–849.
- Cohen, A., Dahmen, W. and DeVore, R. [2009], ‘Compressed sensing and best k-term approximation’, *Journal of the AMS* **22**(1), 211–231.
- Deshpande, Y., Montanari, A. and Richard, E. [2014], Cone-constrained principal component analysis, in ‘Advances in Neural Information Processing Systems’, pp. 2717–2725.
- Donoho, D. and Huo, X. [2001], ‘Uncertainty principles and ideal atomic decomposition’, *IEEE Transactions on Information Theory* **47**(7), 2845–2862.
- Donoho, D. L. and Tanner, J. [2005], ‘Sparse nonnegative solutions of underdetermined linear equations by linear programming’, *Proc. of the National Academy of Sciences* **102**(27), 9446–9451.
- Donoho, D. L. and Tsaig, Y. [2008], ‘Fast solution of ℓ_1 -norm minimization problems when the solution may be sparse-norm minimization problems when the solution may be sparse’, *Information Theory, IEEE Transactions on* **54**(11), 4789–4812.
- Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. [2004], ‘Least angle regression’, *Annals of Statistics* **32**(2), 407–499.
- Epelman, M. and Freund, R. M. [2000], ‘Condition number complexity of an elementary algorithm for computing a reliable solution of a conic linear system’, *Mathematical Programming* **88**(3), 451–485.
- Fercoq, O. and Qu, Z. [2016], ‘Restarting accelerated gradient methods with a rough strong convexity estimate’, *arXiv preprint arXiv:1609.07358*.
- Feuer, A. and Nemirovski, A. [2003], ‘On sparse representation in pairs of bases’, *IEEE Transactions on Information Theory* **49**(6), 1579–1581.
- Freund, R. M. and Vera, J. R. [1999a], ‘Condition-based complexity of convex optimization in conic linear form via the ellipsoid algorithm’, *SIAM Journal on Optimization* **10**(1), 155–176.
- Freund, R. M. and Vera, J. R. [1999b], ‘Some characterizations and properties of the “distance to ill-posedness” and the condition measure of a conic linear system’, *Mathematical Programming* **86**(2), 225–260.
- Freund, R. M. and Vera, J. R. [2003], ‘On the complexity of computing estimates of condition measures of a conic linear system’, *Mathematics of Operations Research* **28**(4), 625–648.
- Giselsson, P. and Boyd, S. [2014], Monotonicity and restart in fast gradient methods, in ‘53rd IEEE Conference on Decision and Control’, IEEE, pp. 5058–5063.
- Grant, M., Boyd, S. and Ye, Y. [2001], ‘CVX: Matlab software for disciplined convex programming’.
- Horn, R. A. and Johnson, C. R. [1990], *Matrix analysis*, Cambridge University Press, Cambridge. Corrected reprint of the 1985 original.
- Journée, M., Nesterov, Y., Richtárik, P. and Sepulchre, R. [2008], ‘Generalized power method for sparse principal component analysis’, *arXiv:0811.4724*.
- Juditsky, A., Karzan, F. K. and Nemirovski, A. [2014], ‘On a unified view of nullspace-type conditions for recoveries associated with general sparsity structures’, *Linear Algebra and its Applications* **441**, 124–151.
- Kashin, B. and Temlyakov, V. [2007], ‘A remark on compressed sensing’, *Mathematical notes* **82**(5), 748–755.
- Luss, R. and Teboulle, M. [2013], ‘Conditional gradient algorithms for rank-one matrix approximations with a sparsity constraint’, *SIAM Review* **55**(1), 65–98.
- Negahban, S., Yu, B., Wainwright, M. J. and Ravikumar, P. K. [2009], A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers, in ‘Advances in Neural Information Processing Systems’, pp. 1348–1356.
- Nemirovskii, A. and Nesterov, Y. E. [1985], ‘Optimal methods of smooth convex minimization’, *USSR Computational Mathematics and Mathematical Physics* **25**(2), 21–30.
- Nesterov, Y. [1983], ‘A method of solving a convex programming problem with convergence rate $O(1/k^2)$ ’, *Soviet Mathematics Doklady* **27**(2), 372–376.
- Nesterov, Y. [2005], ‘Smooth minimization of non-smooth functions’, *Mathematical Programming* **103**(1), 127–152.
- Obozinski, G., Jacob, L. and Vert, J.-P. [2011], ‘Group lasso with overlaps: the latent group lasso approach’, *arXiv preprint arXiv:1110.0413*.
- O’Donoghue, B. and Candes, E. [2015], ‘Adaptive restart for accelerated gradient schemes’, *Foundations of computational mathematics* **15**(3), 715–732.
- Ordóñez, F. and Freund, R. M. [2003], ‘Computational experience and the explanatory value of condition measures for linear optimization’, *SIAM Journal on Optimization* **14**(2), 307–333.

- Oymak, S. and Hassibi, B. [2010], ‘New null space results and recovery thresholds for matrix rank minimization’, *arXiv preprint arXiv:1011.6326*.
- Pajor, A. and Tomczak-Jaegermann, N. [1986], ‘Subspaces of small codimension of finite-dimensional banach spaces’, *Proceedings of the American Mathematical Society* **97**(4), 637–642.
- Peña, J. [2000], ‘Understanding the geometry of infeasible perturbations of a conic linear system’, *SIAM Journal on Optimization* **10**(2), 534–550.
- Recht, B., Xu, W. and Hassibi, B. [2008], Necessary and sufficient conditions for success of the nuclear norm heuristic for rank minimization, in ‘Decision and Control, 2008. CDC 2008. 47th IEEE Conference on’, IEEE, pp. 3065–3070.
- Renegar, J. [1995a], ‘Incorporating condition measures into the complexity theory of linear programming’, *SIAM Journal on Optimization* **5**(3), 506–524.
- Renegar, J. [1995b], ‘Linear programming, complexity theory and elementary functional analysis’, *Mathematical Programming* **70**(1-3), 279–351.
- Renegar, J. [2001], *A mathematical view of interior-point methods in convex optimization*, Vol. 3, Siam.
- Rockafellar, R. T. [1970], *Convex Analysis*, Princeton University Press., Princeton.
- Roulet, V. and d’Aspremont, A. [2017], ‘Sharpness, restart and acceleration’, *arXiv preprint arXiv:1702.03828*.
- Su, W., Boyd, S. and Candes, E. [2014], A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights, in ‘Advances in Neural Information Processing Systems’, pp. 2510–2518.
- Tibshirani, R. [1996], ‘Regression shrinkage and selection via the LASSO’, *Journal of the Royal statistical society, series B* **58**(1), 267–288.
- Van De Geer, S. and Bühlmann, P. [2009], ‘On the conditions used to prove oracle results for the lasso’, *Electronic Journal of Statistics* **3**, 1360–1392.
- Vera, J. C., Rivera, J. C., Peđa, J. and Hui, Y. [2007], ‘A primal–dual symmetric relaxation for homogeneous conic systems’, *Journal of Complexity* **23**(2), 245–261.
- Vershynin, R. [2011], *Lectures in Geometric Functional Analysis*, In preparation.
URL: <http://www-personal.umich.edu/~romanv/papers/GFA-book/GFA-book.pdf>
- Yen, I. E.-H., Hsieh, C.-J., Ravikumar, P. K. and Dhillon, I. S. [2014], Constant nullspace strong convexity and fast convergence of proximal methods under high-dimensional settings, in ‘Advances in Neural Information Processing Systems’, pp. 1008–1016.
- Zhong, Y. and Boumal, N. [2017], ‘Near-optimal bounds for phase synchronization’, *arXiv preprint arXiv:1703.06605*.
- Zhou, Z. and So, A. M.-C. [2017], ‘A unified approach to error bounds for structured convex optimization problems’, *Mathematical Programming* pp. 1–40.
- Zhou, Z., Zhang, Q. and So, A. M.-C. [2015], 11, p-norm regularization: Error bounds and convergence rate analysis of first-order methods, in ‘Proceedings of the 32nd International Conference on Machine Learning, (ICML)’, pp. 1501–1510.

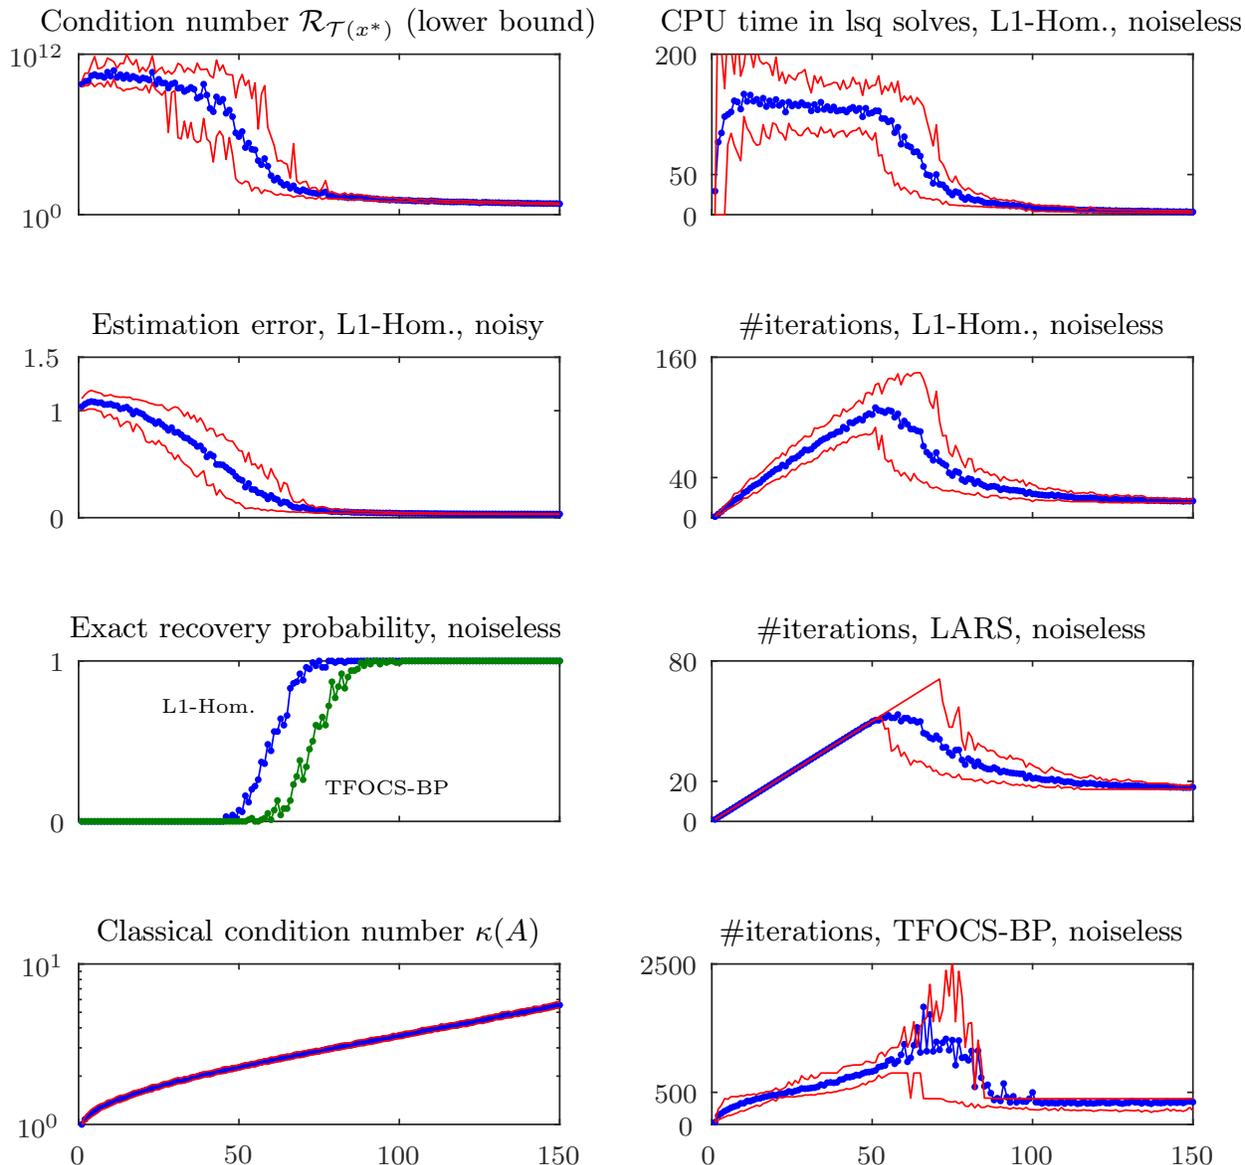


FIGURE 4. We plot the cone-restricted condition number of A (upper left), explaining both the computational complexity of problem (ℓ_1 recovery) (right column) and the statistical complexity of problem (Robust ℓ_1 recovery) (second on the left). Central curves represent the mean (geometric mean in log-scale plots), red curves correspond to 10th and 90th percentile. We observe that high computing times (peaks in the right column) are directly aligned with instances where sparse recovery barely holds/fails (left), i.e. near the phase transition around $n = 70$, where the distance to feasibility for problem $(P_{A,\mathcal{T}(x)})$ also follows a phase transition.

APPENDIX A. PRACTICAL OPTIMAL RESTART SCHEME

In Section 1.2 we quickly give optimal restart schemes in terms of a potentially non-integer clock. Following corollary details optimal scheme for an integer optimal clock.

Corollary A.1. *Given a coding matrix $A \in \mathbb{R}^{n \times p}$ and an original signal $x^* \in \mathbb{R}^p$ such that sharpness bound (Sharp) is satisfied with $\gamma > 0$, running Algorithm *Restart* with ρ^* and $t = \lceil t^* \rceil$ where ρ^* and t^* are defined in (4) ensures that after $K \geq 1$ restarts, i.e. $N = K \lceil t^* \rceil$ total number of iterations,*

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq \exp\left(-\frac{N\gamma}{2e\sqrt{p} + \gamma}\right) \epsilon_0. \quad (27)$$

Proof. Denote $\delta = \lceil t^* \rceil - t^* \in [0, 1[$. As $\lceil t^* \rceil \geq t^*$ (3) is ensured for ρ^* . At the K^{th} restart, $N = K(t^* + \delta)$, and

$$\|\hat{y}\|_1 - \|x^*\|_1 \leq e^{-K} \epsilon_0 = \exp(-N/(t^* + \delta)) \leq \exp(-N/(t^* + 1)).$$

Replacing t^* by its value gives the result. ■

APPENDIX B. REMARK ON SPARSITY INDUCING NORMS

We quickly discuss the framework of Juditsky et al. [2014] for sparsity inducing norms and show that it can be simplified. We first recall the definition.

Definition B.1. (Sparsity structure [Juditsky et al., 2014]) *A sparsity structure on a Euclidean space E is defined as a norm $\|\cdot\|$ on E , together with a family \mathcal{P} of linear maps of E into itself, satisfying three assumptions:*

- (1) *Every $P \in \mathcal{P}$ is a projector, $P^2 = P$,*
- (2) *Every $P \in \mathcal{P}$ is assigned a weight $\nu(P) \geq 0$ and a linear map \bar{P} on E such that $P\bar{P} = 0$,*
- (3) *For any $P \in \mathcal{P}$ and $u, v \in E$, one has*

$$\|P^*u + \bar{P}^*v\|_* \leq \max(\|u\|_*, \|v\|_*),$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$ and P^* is the conjugate mapping of the linear map P .

The last condition in Definition B.1 is arguably the least intuitive and following Lemma connects it with the more intuitive notion of decomposable norm.

Lemma B.2. *Condition (3) above, which reads*

$$\|P^*u + \bar{P}^*v\|_* \leq \max(\|u\|_*, \|v\|_*),$$

for any $u, v \in E$, is equivalent to

$$\|w\| \geq \|Pw\| + \|\bar{P}w\|,$$

for any $w \in E$.

Proof. Denote $f : (u, v) \rightarrow \|P^*u + \bar{P}^*v\|_*$ and $g : (u, v) \rightarrow \max(\|u\|_*, \|v\|_*)$. Since f and g are non-negative, continuous convex functions, $f^2/2$ and $g^2/2$ are also convex continuous and following equivalences hold

$$f \leq g \Leftrightarrow \frac{f^2}{2} \leq \frac{g^2}{2} \Leftrightarrow \left(\frac{f^2}{2}\right)^* \geq \left(\frac{g^2}{2}\right)^*,$$

using that for a convex continuous function h , $h^{**} = h$. Now combining the conjugacy result for squared norm [Boyd and Vandenberghe, 2004, Example 3.27] showing that the conjugate of a squared norm $\|x\|^2/2$ is the squared conjugate norm $\|x\|_*^2/2$, with the result in [Rockafellar, 1970, Th. 16.3], we get

$$\left(\frac{f^2}{2}\right)^*(s, t) = \inf_w \{\|w\|^2/2 : Pw = s, \bar{P}w = t\},$$

where the infimum is $+\infty$ if the constraints are infeasible. Then the dual of the norm g is $(s, t) \rightarrow \|s\| + \|t\|$ therefore condition (3) is equivalent to

$$\inf_w \{\|w\| : Pw = s, \bar{P}w = t\} \geq \|s\| + \|t\|,$$

for any $s, t \in E$, which reads

$$\|w\| \geq \|Pw\| + \|\bar{P}w\|,$$

for any $w \in E$. ■

DI, ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.

INRIA SIERRA TEAM.

E-mail address: `vincent.roulet@inria.fr`

MATHEMATICS DEPARTMENT,

PRINCETON UNIVERSITY, PRINCETON NJ 08544, USA.

E-mail address: `nboumal@math.princeton.edu`

CNRS & DI, ÉCOLE NORMALE SUPÉRIEURE, PARIS, FRANCE.

E-mail address: `aspremon@di.ens.fr`