

Sharpness and well-conditioning of nonsmooth convex formulations in statistical signal recovery

Lijun Ding ^{*} Alex L. Wang [†]

July 14, 2023

Abstract

We study a sample complexity vs. conditioning tradeoff in modern signal recovery problems where convex optimization problems are built from sampled observations. We begin by introducing a set of condition numbers related to sharpness in ℓ_p or Schatten- p norms ($p \in [1, 2]$) based on nonsmooth reformulations of a class of convex optimization problems, including sparse recovery, low-rank matrix sensing, covariance estimation, and (abstract) phase retrieval. In each of the recovery tasks, we show that the condition numbers become dimension independent constants once the sample size exceeds some constant multiple of the recovery threshold. Structurally, this result ensures that the inaccuracy in the recovered signal due to both observation noise and optimization error is well-controlled. Algorithmically, such a result ensures that a new first-order method for solving the class of sharp convex functions in a given ℓ_p or Schatten- p norm, when applied to the nonsmooth formulations, achieves nearly-dimension-independent linear convergence.

1 Introduction

This paper studies a *sample complexity vs. convex optimization conditioning* tradeoff in modern signal recovery problems such as abstract phase retrieval (henceforth phase retrieval) [14], sparse recovery [15], low-rank matrix sensing [9, 50], and covariance estimation [21].

Statistical signal recovery via convex optimization. To set the stage, consider recovering an unknown element of a vector space, $x^\natural \in V$, called the signal, from a linear sensing operator, $\mathcal{A} : V \rightarrow \mathbb{R}^m$, and a vector of measurements, $b = \mathcal{A}(x^\natural)$. In short, we need to

recover $x^\natural \in V$ from (b, \mathcal{A}) where $b = \mathcal{A}(x^\natural) \in \mathbb{R}^m$.

In a standard statistical signal recovery setting, each coordinate of the linear map \mathcal{A} is sampled from a fixed distribution and m is the number of samples, i.e, $\mathcal{A}^*(e_1), \dots, \mathcal{A}^*(e_m)$ are i.i.d, where \mathcal{A}^* is the adjoint map and $e_i, i = 1, \dots, m$, are standard basis vectors. The convex optimization approach to this problem, pioneered in [15, 50], begins by constructing a convex optimization problem

$$\min_{x \in V} \left\{ f(x) : \begin{array}{l} \mathcal{A}(x) = b \\ x \in \mathcal{K} \end{array} \right\}. \quad (1)$$

^{*}Wisconsin Institute for Discovery, University of Wisconsin-Madison. Madison, WI, USA(liding47@wisc.edu)

[†]Centrum Wiskunde & Informatica and Purdue University. West Lafayette, IN, USA(wang5984@purdue.edu)

Here, the choice of the convex objective f and the convex set \mathcal{K} reflect prior knowledge about the signal x^\natural . For example, we may set f to be the vector ℓ_1 -norm if x^\natural is sparse [15] or \mathcal{K} to be the positive semidefinite (PSD) cone when x^\natural is PSD¹ [14]. A tremendous effort in the past [9, 14–17, 20, 23, 49, 50] showed that, once the sample size m exceeds a small multiple of the information-theoretic threshold², the convex optimization problem admits x^\natural as its *unique* minimizer with high probability (w.h.p.).³

Issues and challenges. Unfortunately, convexity and uniqueness of solution are not the end of the story as convex optimization problems themselves may not be well-conditioned. In these situations, the recovery process may suffer from

- *measurement error*: when the observation vector b is corrupted by noise, the optimal solution of (1) may deviate greatly from x^\natural or may even cease to exist;
- *optimization error*: the problem (1) may admit near-optimal and near-feasible solutions in terms of function value and constraint violation that are far from the true signal x^\natural ;
- *slow algorithm convergence*: the iteration complexity of existing first-order algorithms for solving (1) may depend polynomially on the dimension due to polynomially poor conditioning of (1).⁴ Such dependence weakens the applicability of existing first-order methods on large-scale instances of (1).

Our contribution. In this paper, we introduce a notion of conditioning related to *sharpness* for (1) based on an *unconstrained nonsmooth reformulation*⁵ of (1) that allows us to quantitatively control each of the above issues. Our contributions are summarized below:

- Statistically, we prove that a variety of problems from statistical signal recovery are well-conditioned in terms of ℓ_1 or *Schatten-1 norm* under standard statistical assumptions once m is taken to be a constant factor above the information-theoretic threshold (Section 4). Numerically, we observe that the convex optimization problems encountered in these signal recovery problems are ill-conditioned near the known thresholds for successful recovery [2, 17]. Furthermore, (1) is *not* well-conditioned in standard Euclidean notions of conditioning leading to poor numerical performance of ℓ_2 based algorithms such as the subgradient method (Section 8).
- Structurally, we show that if (1) is better conditioned, near-optimal and near-feasible solutions are closer to the true signal, and noise/error in the observations has less of an impact on the signal recovered via the nonsmooth reformulations (Sections 2 and 5). Furthermore, our notion of conditioning implies an important regularity condition known as strict complementarity (Section 7).
- Algorithmically, we develop a first-order method for the broader problem class of unconstrained nonsmooth convex problems that are well-conditioned in terms of ℓ_p or Schatten- p norms

¹A matrix is PSD if it is symmetric and all its eigenvalues are nonnegative.

²This threshold is usually proportional to the dimension of the manifold or the union of manifolds where x^\natural belongs, e.g., the fixed rank manifold or the set of sparse vectors.

³An event happens w.h.p. if the probability is larger than $1 - \exp(-cm)$ for some numerical constant c .

⁴These polynomial factors may be even worse when measuring error in terms of distance to the true solution as opposed to function value.

⁵See Section 2 or the following subsection for a description of the reformulation and the sharpness in terms of a particular norm.

$p \in [1, 2]$. This algorithm has an iteration complexity that is nearly dimension-independent⁶ (Section 6). Numerically, we observe its iteration complexity stays nearly constant for (1) when the dimension increases. On the other hand, methods that are optimal for well-conditioned problems in terms of Euclidean norm, e.g., the subgradient method with Polyak step sizes, deteriorate with dimension (Section 8).

We illustrate our results more concretely in the setting of phase retrieval in the following subsection.

1.1 Vignette: phase retrieval

The phase retrieval problem asks us to recover an unknown rank-one PSD matrix $X^\natural \in \mathbb{S}^n$ given a vector of m measurements $b = \mathcal{A}(X^\natural)$ where $\mathcal{A} : \mathbb{S}^n \rightarrow \mathbb{R}^m$ is defined by $\mathcal{A}(X)_i = g_i^\top X g_i$ for i.i.d. $g_i \sim N(0, I_n/m)$. In other words, $\mathcal{A}^*(e_i) = g_i g_i^\top$ for all $i = 1, \dots, m$. Here and throughout, \mathbb{S}^n and \mathbb{S}_+^n denote the vector space of symmetric matrices and the cone of PSD matrices respectively. In [10], Candès and Li established an *exact recovery* result: If $m = \Theta(n)$, then, w.h.p., X^\natural is the unique optimizer for

$$\min_{X \in \mathbb{S}^n} \left\{ \text{tr}(X) : \begin{array}{l} \mathcal{A}(X) = b \\ X \in \mathbb{S}_+^n \end{array} \right\}. \quad (2)$$

Nonsmooth reformulation and conditioning. To measure the conditioning of (2), we introduce the following penalized nonsmooth version of (2):

$$\min_{X \in \mathbb{S}^n} F_{r,\ell}(X) := \text{tr}(X) + r \|\mathcal{A}(X) - b\|_1 + \ell \text{dist}_{\|\cdot\|_1}(X, \mathbb{S}_+^n). \quad (3)$$

Here, the distance $\text{dist}_{\|\cdot\|_1}(X, \mathbb{S}_+^n) := \inf_{Y \in \mathbb{S}_+^n} \|Y - X\|_1$ is measured in the Schatten-1 norm $\|\cdot\|_1$, i.e., the sum of the singular values, and $r, \ell \geq 0$ are penalty parameters on the violations of $\mathcal{A}(X) = b$ and $X \in \mathbb{S}_+^n$ respectively.

We emphasize that we penalize the violation of $\mathcal{A}(X) = b$ by $\|\mathcal{A}(X) - b\|_1$ and not the more common lasso-type error $\|\mathcal{A}(X) - b\|_2^2$ and we use the Schatten-1 norm, $\|\cdot\|_1$, in the distance function. These choices lead to a nonsmooth problem even when restricted to the interior of the PSD cone. On the other hand, such choices open the possibility for the function (3) to be μ -sharp around X^\natural w.r.t. $\|\cdot\|_1$ for some $\mu > 0$:

$$\underbrace{\text{tr}(X) + r \|\mathcal{A}(X) - b\|_1 + \ell \text{dist}_{\|\cdot\|_1}(X, \mathbb{S}_+^n)}_{F_{r,\ell}(X)} - \underbrace{\text{tr}(X^\natural)}_{F_{r,\ell}(X^\natural)} \geq \mu \|X - X^\natural\|_1, \quad \text{for all } X \in \mathbb{S}^n. \quad (4)$$

Note that such a bound necessitates that X^\natural is the unique minimizer of both (1) and (3). When such a bound holds, we will think of the parameters (μ, r, ℓ) as partially describing the conditioning of (1). Finally, in order to make our notion of conditioning “aware” of scaling, we additionally track the Lipschitz constant, L , of (3) with respect to $\|\cdot\|_1$:

$$|F_{r,\ell}(X) - F_{r,\ell}(Y)| \leq L \|X - Y\|_1, \quad \text{for all } X, Y \in \mathbb{S}^n.$$

We use the set of numbers (μ, r, ℓ, L) to measure the conditioning of (1). The sharpness parameter μ in some sense will control absolute notions of error, while the condition number $\kappa := L/\mu$ will control relative notions of error and convergence rates of first-order methods.⁷

⁶Its per-iteration complexity mainly requires one call of \mathcal{A} and \mathcal{A}^* and an additional singular value or eigenvalue decomposition in the matrix case. This could be significantly smaller than other first order methods that need to solve a linear system, e.g., ADMM.

⁷One may replace the Schatten-1 norm with other norms in these definitions. The benefit of the Schatten-1 norm, as we will discuss, is that the conditioning in our applications can be shown to be *dimension-independent*.

Our results for phase retrieval. With the above notions and notations, let us describe our results more concretely.

- *Well-conditioning:* By adapting well-known proofs for exact recovery [15, 50], we show that (3) is μ -sharp and L -Lipschitz with $\mu, L = \Theta(1)$ for $r, \ell = \Theta(1)$ w.h.p. once $m = \Theta(n)$ (See Theorem 1). Note this implies the reformulation (3) is indeed exact.
- *Optimization error:* In practice, we are unable to solve (2) exactly and must resort to numerical optimization. Using the μ -sharpness in (4), any ϵ -suboptimal solution to (2), or more generally to (3), satisfies $\|X - X^\natural\|_* \leq \epsilon/\mu$.
- *Measurement error:* Suppose that instead of observing $b = \mathcal{A}(X^\natural)$, we observe $\tilde{b} = \mathcal{A}(X^\natural) + \delta$. It is known that (2) actually admits a unique *feasible* solution X^\natural w.h.p. [10], thus (2) with b replaced by \tilde{b} is likely to be infeasible. On the other hand, given any optimizer \tilde{X} of (3) with \tilde{b} in place of b , we may use sharpness to bound $\|\tilde{X} - X^\natural\|_1 \leq O(\|\delta\|_1/\mu)$ (see Proposition 2).

In a similar vein, we may consider a setting where b is corrupted by a sparse vector δ . Sharpness tells us that the optimizer of (3) is *unchanged* if $|\text{supp}(\delta)|/m = O(\mu)$, i.e., if up to an $O(\mu)$ fraction of the entries of b are corrupted (see Proposition 3 and Example 1).

- *Strict complementarity:* Strict complementarity is an important regularity condition in structural optimization [27, 62]. We show that sharpness implies strict complementarity so that strict complementarity holds for (2). This is surprising as (2) *does not* even satisfy Slater’s condition due to having a unique feasible solution [10].
- *Algorithms:* To solve large-scale instances of (2), we develop a restarted mirror descent algorithm for *general* convex functions that are μ -sharp and L -Lipschitz in terms of a Schatten- p norm with $p \in [1, 2]$ (see Section 6). This algorithm is capable of producing an ϵ -suboptimal solution to (3) in

$$O\left(\kappa^2 \log(n) \log\left(\frac{1}{\epsilon}\right)\right)$$

iterations, i.e., with nearly-dimension-independent linear convergence. Furthermore, these convergence guarantees hold after appropriate modifications in the presence of corruption or noise (see Propositions 2 and 3) with guarantees depending on κ .

We note that our results on optimization error, measurement error, strict complementarity, and algorithms above depend solely on the deterministic assumption of sharpness and “well-conditioning.” The statistical assumptions are only used to prove that the conditioning (μ, r, ℓ, L) is well-behaved w.h.p.

1.2 Related work

To better position our work in the literature, in this section, we discuss related work on error bounds and well-conditioning, sharpness and well-conditioning in statistical recovery, and first-order methods for minimizing sharp functions.

Establishing error bounds and well-conditioning. The sharpness condition in (4) can be viewed as an error bound for (1). Error bounds and conditioning are central to optimization

problems both structurally and algorithmically [1, 27, 39, 40, 43, 62]. Consider the following general error bound for (1):

$$f(x) + r \|A(x) - b\|_1 + \ell \operatorname{dist}_{\|\cdot\|}(x, \mathcal{K}) - f(x^\dagger) \geq \mu \|x - x^\dagger\|^h, \quad \text{for all } x \in V,$$

where, in addition to the r and ℓ introduced in (4), we also have the power h , and a general norm $\|\cdot\|$. Most methods for establishing error bounds either require that (1) is linear [38] or that regularity conditions such as Slater’s condition [6, 61], strict complementarity [24, 56], or nondegeneracy conditions [42] hold. Additional work has studied error bounds in generic semialgebraic settings [26]. Results of this type generally place an emphasis on establishing the power of h (usually to 1 or 2) and provide only weak bounds on μ , r , and ℓ . This is natural as the power h determines the linear or sublinear convergence rate [3, 39]. On the other hand, if μ , r , or ℓ depend polynomially on the dimension, then first-order algorithms may still be doomed to poor or dimension-dependent convergence rates. This work connects the quantities μ , r , and ℓ to the quantitatively better-understood restricted isometry property (RIP). This connection gives precise estimates of (μ, r, ℓ, L) in our settings and shows that these quantities are in fact dimension independent.

Sharpness and well conditioning in statistical recovery. A series of works [18, 19, 29, 41] studied the local landscape of *nonconvex* nonsmooth formulations in statistical recovery near the signal, particularly in low-rank matrix recovery. They showed that near the ground truth X^\dagger , the natural loss function is sharp, and the local condition number (in terms of the Frobenius or Schatten-2 norm) does not depend on the dimension directly. Hence, they can apply off-the-shelf algorithms, such as subgradient and prox-linear methods [22, 28] to achieve quick convergence to X^\dagger . Note that results of this type are local, and the algorithms require careful initialization to achieve provable guarantees. Moreover, because the matrix space in the nonconvex formulation is actually low-dimensional, the difference between the Schatten- p norm for $p = 1, 2$ is insignificant. This fact results in simpler analyses of conditioning and algorithmic convergence due to the unique properties of inner products and their induced norms. Another subtlety is that the sharpness parameter in the nonconvex setting depends on the condition number of the signal matrix X^\dagger in addition to the RIP parameters. In contrast, the results here only depend on the RIP parameters.

First-order methods for minimizing sharp functions. Early work in this area established linear convergence of the subgradient method (with appropriate step sizes) for μ -sharp L -Lipschitz functions in the Euclidean norm [32, 35, 48]. These methods and their proofs are adapted to the Euclidean norm and incur a *polynomial dependence on the dimension* when applied to μ -sharp L -Lipschitz functions in ℓ_1 or Schatten-1 norms.

Similar guarantees can be derived as a consequence of restarting schemes. Perhaps the earliest work discussing restarting schemes is [44], where a restarting scheme was developed for convex Hölder-smooth minimization in ℓ_p spaces. More recent work [43, 53, 54, 59] has used restarting schemes to accelerate variants of (sub)gradient descent in the presence of different growth conditions, e.g., a local guarantee of the form⁸

$$f(x) - \min_x f(x) \geq \mu \operatorname{dist}(x, \mathcal{X})^h$$

for all x close enough to the set of minimizers, \mathcal{X} . Here, $\operatorname{dist}(\cdot, \cdot)$ is typically measured in the Euclidean norm and $h \in [1, \infty)$ captures the Hölderian growth of f . Taking $h = 1$ recovers

⁸Other forms of the necessary “growth” are possible in different settings.

the usual definition of sharpness. In this setting, restarted subgradient descent achieves a linear convergence rate of $O(\kappa^2 \log(1/\epsilon))$ [59]. Adaptive versions of these restart schemes have also been developed [51, 54] that do not require the sharpness parameter μ to be specified but incur an additional $\log(1/\epsilon)$ factor in the convergence rate.

We emphasize that much of the preceding literature on first-order methods for sharp convex functions focuses on the Euclidean case. The only exception we know of is the work of [54], which develops a restarted mirror descent method based on a property inspired by sharpness in the Euclidean setting. Roulet and d’Aspremont [54] show that their method can be used to achieve $O(\kappa^2 \log(n) \log(1/\epsilon))$ convergence on this particular function class. Unfortunately, their proxy for sharpness is quite restrictive and does not hold in settings where sparse or low rank solutions are expected. Our work in Section 6 is closely related to [54] but uses a different specification of the restarting mechanism. This change allows us to use a more natural definition of sharpness that is also much broader (see Remark 7 for a thorough comparison).

1.3 Outline and notation

Outline. We begin, in Section 2, by setting up notation and defining sharpness for problems of the form (1). We will quantify sharpness and conditioning by a set of parameters (μ, r, ℓ, L) . In Sections 3 and 4, we establish well-conditioning and sharpness under the Gaussian sensing models for the problems of sparse recovery, low-rank matrix and bilinear sensing, and covariance estimation in the absence of noise. Our proof relies on the well-known restricted isometry property (RIP). In Section 5, we investigate how sharp problems of the form (1) behave in the presence of noise. In Section 6, we describe and analyze a variant of restarted mirror descent (RMD) that has linear and nearly dimension-independent convergence rates when applied to sharp functions in ℓ_p or Schatten- p norms (see Assumption 1). In conjunction with the results of Section 5, we may apply RMD to obtain linear nearly dimension-independent convergence rates even in the presence of noise for those statistical recovery problems. In Section 7, we show that sharpness in (1) implies a common regularity known as strict complementarity (see (15) and (16)). Finally, in Section 8 we implement a Polyak-variant of RMD and experimentally test its convergence rate as the number of samples m and the dimension of the problem is varied in statistical recovery problems.

A note on ℓ_p and Schatten- p norms. Let $p \in [1, \infty]$. We will overload notation and use $\|x\|_p$ for the ℓ_p norm when $x \in \mathbb{R}^n$ and the Schatten- p norm when $x \in \mathbb{R}^{n \times N}$. Recall that the nuclear norm is equivalent to the Schatten-1 norm, the Frobenius norm is equivalent to the Schatten-2 norm, and the spectral norm is equivalent to the Schatten- ∞ norm. To streamline the text, we will also take “ ℓ_p norm” of a matrix to mean its Schatten- p norm.

2 Preliminaries on sharpness and conditioning

Given two finite-dimensional normed real vector spaces V and W , we consider the following abstract problem, a slight generalization of (1),

$$\min_{x \in V} \left\{ f(x) : \begin{array}{l} \mathcal{A}(x) = b \\ x \in \mathcal{K} \end{array} \right\}. \quad (\text{P})$$

The function $f : V \rightarrow \mathbb{R}$ is convex, the map $\mathcal{A} : V \rightarrow W$ is linear, and $\mathcal{K} \subseteq V$ is a closed convex cone (possibly all of V). Each vector space V and W is equipped with a norm $\|\cdot\|$ which may be

different from the Euclidean norm. For notational simplicity, we write all norms as $\|\cdot\|$ with the understanding that the norm is the norm associated with the space of the argument.

In this section, we first formally define the sharpness and conditioning of **(P)**. We then give a dual characterization of sharpness. The dual characterization is useful in showing (i) the exact penalization formulation (5) is robust to noisy b (cf. Section 5), and (ii) sharpness implies a well-known regularity assumption of strict complementarity (cf. Section 7).

2.1 Definition of sharpness and conditioning

For $x \in V$ and a closed nonempty set $\mathcal{X} \subseteq V$, let

$$\text{dist}(x, \mathcal{X}) := \min_{\bar{x} \in \mathcal{X}} \|\bar{x} - x\|.$$

Again, the norm $\|\bar{x} - x\|$ is measured in the space of the argument. In this case, the expression $\bar{x} - x \in V$ is measured in the norm of V . The following definition of sharpness for a convex function (or unconstrained minimization problem) is also referred to as weak sharpness in the literature [33].

Definition 1 (Sharpness). Let V be a finite-dimensional normed real vector space. Let $f : V \rightarrow \mathbb{R}$ be a convex function, $\mathcal{X} \subseteq V$ a nonempty convex set, and let $\mu > 0$. We say that f is μ -sharp around \mathcal{X} if $\mathcal{X} = \arg \min f$ and

$$f(\bar{x}) - \min_{x \in V} f(x) \geq \mu \text{dist}(\bar{x}, \mathcal{X}), \quad \forall \bar{x} \in V. \quad \square$$

We extend the notion of sharpness to a problem of the form **(P)** as follows.

Definition 2 (Sharpness of **(P)**). Consider a problem of the form **(P)**. Let $x^\natural \in V$, $\mu > 0$, and $r, \ell \geq 0$. We say that **(P)** is (μ, r, ℓ) sharp around x^\natural if x^\natural is feasible in **(P)** and

$$F_{r,\ell}(x) := f(x) + r \|\mathcal{A}(x) - b\| + \ell \text{dist}(x, \mathcal{K}) \quad (5)$$

is μ -sharp around x^\natural . Equivalently, if x^\natural is feasible in **(P)** and

$$\|x - x^\natural\| \leq \frac{1}{\mu} \left(f(x) - f(x^\natural) + r \|\mathcal{A}(x) - b\| + \ell \text{dist}(x, \mathcal{K}) \right), \quad \forall x \in V. \quad (6)$$

If $\mathcal{K} = V$, then ℓ is inconsequential, so we will simply say **(P)** is (μ, r) sharp around x^\natural . \square

Remark 1. Suppose **(P)** is (μ, r, ℓ) -sharp around x^\natural for some $\mu > 0$, $r, \ell \geq 0$. Note that x^\natural is necessarily the unique optimal solution of **(P)**. Moreover, we have an error bound: If we numerically solve **(P)** and produce $\tilde{x} \in V$, then we may bound the distance $\|\tilde{x} - x^\natural\|$ according to (6). This bound holds even if $f(\tilde{x}) < f(x^\natural)$, which could happen as \tilde{x} is not necessarily feasible for **(P)**. \square

Recall the standard definition of the Lipschitz constant of a function.

Definition 3 (Lipschitz constant). We say that a convex function $f : V \rightarrow \mathbb{R}$ is L -Lipschitz if

$$|f(x) - f(y)| \leq L \|x - y\|, \quad \forall x, y \in V. \quad \square$$

We are now ready to define the conditioning of **(P)**. Note that the following definition depends on the norms associated with V and W .

Definition 4 (Conditioning and condition number). Consider a problem of the form (P) and suppose $x^\natural \in V$ is its unique optimizer. Suppose for some $r, \ell \geq 0$ that $F_{r,\ell}$ is μ -sharp around x^\natural and L -Lipschitz, with $\mu, L > 0$. The conditioning of (P) is measured by (μ, r, ℓ, L) and the condition number is $\kappa = \frac{L}{\mu}$. \square

2.2 Dual perspective of sharpness

In this section, we apply known characterizations of sharpness from a dual perspective to (P). These results will be useful in the remainder of the paper.

Dual space, norm, linear map, and cone. We start by recalling some basic definitions. Let V^* denote the space of linear functional on V and let the primal-dual pairing $\langle g, x \rangle := g(x)$ for $g \in V^*$ and $x \in V$. We define the norm on V^* as $\|g\| := \max_{x \in V} \{\langle g, x \rangle : \|x\| \leq 1\}$. For example, if $V = (\mathbb{R}^n, \|\cdot\|_1)$ is the space of n -dimensional vectors with the ℓ_1 -norm, then we have $V^* = (\mathbb{R}^n, \|\cdot\|_\infty)$, where we have identified linear functions on \mathbb{R}^n with vectors in \mathbb{R}^n via the canonical inner product. Given a closed convex cone $\mathcal{K} \subset V$, its dual cone is defined as $\mathcal{K}^* := \{g \in V^* : \langle g, v \rangle \geq 0, \text{ for all } v \in \mathcal{K}\}$. For a linear operator $\mathcal{A} : V \rightarrow W$, the dual linear map of \mathcal{A} is the unique linear operator $\mathcal{A}^* : W^* \rightarrow V^*$ such that⁹

$$\langle w, \mathcal{A}v \rangle = \langle \mathcal{A}^*w, v \rangle \quad \forall v \in V, w \in W^*.$$

The following lemma presents a dual view of μ -sharpness, i.e., in terms of subgradients. This lemma is useful when noise is present in b as shown in Section 5. For any vector space V with norm $\|\cdot\|$, we use $B_V(0, r) := \{x \in V : \|x\| \leq r\}$ to denote the closed ball of radius r centered at $0 \in V$. Given $x \in V$, let $x^\perp := \{g \in V^* : \langle g, x \rangle = 0\}$.

Lemma 1. *Suppose x^\natural is feasible in (P), $\mu > 0$, and $r, \ell \geq 0$. Then, (P) is (μ, r, ℓ) sharp around x^\natural if and only if*

$$B_{V^*}(0, \mu) \subseteq \partial f(x^\natural) - \mathcal{A}^*(B_{W^*}(0, r)) - (\mathcal{K}^* \cap (x^\natural)^\perp \cap B_{V^*}(0, \ell)).$$

Proof. By definition, (P) is (μ, r, ℓ) -sharp around x^\natural iff $F_{r,\ell}(x)$ is μ -sharp around x^\natural iff

$$B_{V^*}(0, \mu) \subseteq \partial F_{r,\ell}(x^\natural).$$

The last equivalence follows from [60] (specifically, the equivalence of properties (i) and (x) in [60, Theorem 3.10.1]). The claim follows then from additivity of the subdifferential and the identities (see [7, Proposition 18.22])

$$\begin{aligned} \partial \left(r \left\| \mathcal{A}(x^\natural) - b \right\| \right) &= -\mathcal{A}^*(B_{W^*}(0, r)), \text{ and} \\ \partial \left(\ell \text{dist}(x^\natural, \mathcal{K}) \right) &= -(\mathcal{K}^* \cap (x^\natural)^\perp \cap B_{V^*}(0, \ell)). \quad \blacksquare \end{aligned}$$

The following lemma characterizes whether (1) is (μ, r, ℓ) sharp around x^\natural for *some* $\mu > 0$ and $r, \ell \geq 0$. This lemma will be useful in proving negative results in Section 7. We defer the proof to Appendix A.

⁹If both spaces V and W are equipped with some inner products $\langle \cdot, \cdot \rangle_V$ and $\langle \cdot, \cdot \rangle_W$, then \mathcal{A}^* is also the adjoint map $\bar{\mathcal{A}} : W \rightarrow V$ defined via $\langle w, \mathcal{A}v \rangle_W = \langle \bar{\mathcal{A}}w, v \rangle_V$ if we use the inner product $\langle \cdot, \cdot \rangle_V$ to identify V and V^* , and use $\langle \cdot, \cdot \rangle_W$ to identify W and W^* .

Lemma 2. Suppose x^\natural is feasible in (P). Then, there exists $\mu > 0$, and $r, \ell \geq 0$ such that (P) is (μ, r, ℓ) sharp around x^\natural if and only if

$$0 \in \text{int} \left(\partial f(x^\natural) - \text{range}(\mathcal{A}^*) - \left(\mathcal{K}^* \cap (x^\natural)^\perp \right) \right). \quad (7)$$

Remark 2. Note that (7) does not depend on the norms of V or W . In other words, whether (P) is (μ, r, ℓ) sharp around x^\natural for some $\mu > 0$, $r, \ell \geq 0$ is a property of the underlying vector spaces, whereas the values of (μ, r, ℓ) are a property of the normed spaces, i.e., depends on the choices of the norms. \square

3 RIP-based sharpness and conditioning in signal recovery

This section shows that under bounds on restricted isometry (defined below), sharpness holds for three archetypal modern signal recovery problems: sparse vector recovery, low rank matrix recovery, and covariance estimation. Phase retrieval can also be thought of as a special case of covariance estimation. The signal to be recovered, the optimization formulation in (P), and the choices of norms are the following.

Definition 5. We consider the following signal recovery problems that can be formulated as (P) and the corresponding choices for V , f , and \mathcal{K} .

- Sparse vector recovery: recover a k -sparse vector $x^\natural \in V = \mathbb{R}^n$. In the formulation (P), the objective $f(x) = \|x\|_1$, the set $\mathcal{K} = \mathbb{R}^n$, and the input space $V = (\mathbb{R}^n, \|\cdot\|_1)$.
- Low rank matrix recovery : recover a rank k matrix $X^\natural \in V = \mathbb{R}^{n \times N}$. In the formulation (P), the objective $f(X) = \|X\|_1$, the set $\mathcal{K} = \mathbb{R}^n$, and the input space $V = (\mathbb{R}^{n \times N}, \|\cdot\|_1)$.
- Covariance estimation: recover a rank k PSD matrix $X^\natural \in V = \mathbb{S}^n$. In the formulation (P), the objective is the trace $f(X) = \text{tr}(X)$, the set \mathcal{K} is the set of PSD matrices \mathbb{S}_+^n , and the input space $V = (\mathbb{S}^n, \|\cdot\|_1)$. \square

The space W and its norm are, as yet, unspecified. We have some freedom in the choice of the norm as long as crucial bounds on the restricted isometry property (RIP) hold in the norm on W .

Definition 6 (Restricted isometry property). Let k' be a positive integer and $\mathcal{A} : V \rightarrow W$ be a linear operator. We will let $\text{RIP}_{k'}^-(\mathcal{A})$ and $\text{RIP}_{k'}^+(\mathcal{A})$ denote any valid *uniform* lower bound and upper bound on the quantity $\|\mathcal{A}x\| / \|x\|_2$ as $x \in V$ ranges over all elements in V with support size or rank at most k' . \square

Our main result of this section is the following:

Proposition 1. Consider one of the three problems defined in Definition 5. Let $c = 1$ for sparse vector recovery and $c = 2$ for the other two cases. Let $k' > 0$ and $\epsilon > 0$ and suppose $\text{RIP}_{k'}^+(\mathcal{A}) \geq 1$ and

$$\sqrt{\frac{k'}{ck}} \left(\frac{\text{RIP}_{ck+k'}^-(\mathcal{A})}{\text{RIP}_{k'}^+(\mathcal{A})} \right) \geq 1 + \epsilon.$$

Then, (P) is $\left(\frac{\epsilon}{2+\epsilon}, \sqrt{k'}, 2 \right)$ sharp around x^\natural . Furthermore, the Lipschitz constant L is no more than $3 + \sqrt{k'} \text{RIP}_1^+(\mathcal{A})$.

To put this result in context, we will see in Section 4 that the premise of Proposition 1 holds once the sample size m is a small multiple of the recovery threshold w.h.p. with $\epsilon = 2$ and $k' = O(k)$, and $\text{RIP}_1^+(\mathcal{A}) = O(1)$. Hence, Problem (P) is $\left(\frac{1}{2}, O(\sqrt{k}), 2\right)$ sharp around x^\natural with a Lipschitz constant no more than $O(\sqrt{k})$ w.h.p.

The rest of the section consists of proofs of Proposition 1 for the three different settings. We start with the proof of sharpness of matrix sensing and discuss how this proof can be modified for the other two settings. We then bound the Lipschitz constant for all three settings.

3.1 Proof of sharpness in matrix sensing

This section proves Proposition 1 for the setting of low-rank matrix sensing. Let $\Delta \in \mathbb{R}^{n \times N}$ be arbitrary. Our goal is to show that

$$\left\|X^\natural + \Delta\right\|_1 + \sqrt{k'} \|\mathcal{A}(\Delta)\| - \left\|X^\natural\right\|_1 \geq \frac{\epsilon}{2 + \epsilon} \|\Delta\|_1. \quad (8)$$

Change of basis. Without loss of generality, we may work in a basis such that

$$X^\natural = \begin{pmatrix} X_{1,1}^\natural & 0_{k \times (N-k)} \\ 0_{(n-k) \times k} & 0_{(n-k) \times (N-k)} \end{pmatrix} \quad \text{and} \quad \Delta = \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & \Delta_{2,2} \end{pmatrix},$$

where $X_{1,1}^\natural$ is a $k \times k$ diagonal matrix and $\Delta_{2,2} \in \mathbb{R}^{(n-k) \times (N-k)}$ is diagonal with diagonal entries that are nonincreasing in magnitude.

Partitioning of the difference Δ . Let $\delta_{k^\perp} := \text{diag}(\Delta_{2,2})$. The rest of the proof is based on the idea in [15, 50]. We decompose $\delta_{k^\perp} = \sigma_1 + \dots + \sigma_t$ where each $\sigma_i \in \mathbb{R}^{\min(n,N)-k}$. Specifically, let σ_1 extract the first k' coordinates of δ_{k^\perp} , let each subsequent σ_i extract the next k' coordinates of δ_{k^\perp} . Finally, σ_t may have fewer than k' coordinates of δ_{k^\perp} . Let Σ_i denote the matrix of size $n \times N$ with $\text{Diag}(\sigma_i)$ in its bottom-right $(n-k) \times (N-k)$ block or the block itself.

Lower bound on $\|\Delta_{2,2}\|_1$. Since σ_i are disjoint, we can lower bound $\|\Delta_{2,2}\|_1 = \|\delta_{k^\perp}\|_1$ by

$$\|\Delta_{2,2}\|_1 = \|\delta_{k^\perp}\|_1 \geq \sum_{i=1}^{t-1} \|\sigma_i\|_1 = k' \sum_{i=1}^{t-1} \frac{\|\sigma_i\|_1}{k'}.$$

Using δ_{k^\perp} is nonincreasing and each σ_i is k' sparse, we further have

$$\|\Delta_{2,2}\|_1 \geq k' \sum_{i=2}^t \|\sigma_i\|_\infty \geq \sqrt{k'} \sum_{i=2}^t \|\sigma_i\|_2.$$

Next, we use the RIP to obtain

$$\|\Delta_{2,2}\|_1 \geq \frac{\sqrt{k'}}{\text{RIP}_{k'}^+} \sum_{i=2}^t \|\mathcal{A}(\Sigma_i)\| \stackrel{(a)}{\geq} \frac{\sqrt{k'}}{\text{RIP}_{k'}^+} \left(\left\| \mathcal{A} \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,2} & 0 \end{pmatrix} + \mathcal{A}(\Sigma_1) \right\| - \|\mathcal{A}(\Delta)\| \right),$$

where (a) is due to triangle inequality. Using the RIP condition again and $\text{RIP}_{k'}^+ \geq 1$, we have

$$\begin{aligned} \|\Delta_{2,2}\|_1 &\geq \sqrt{k'} \frac{\text{RIP}_{2k+k'}^-}{\text{RIP}_{k'}^+} \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & \Sigma_1 \end{pmatrix} \right\|_2 - \sqrt{k'} \|\mathcal{A}(\Delta)\| \\ &\geq \sqrt{k'} \frac{\text{RIP}_{2k+k'}^-}{\text{RIP}_{k'}^+} \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_2 - \sqrt{k'} \|\mathcal{A}(\Delta)\|. \end{aligned}$$

Lastly, since $\begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix}$ has rank no more than $2k$, we have

$$\begin{aligned} \|\Delta_{2,2}\|_1 &\geq \sqrt{\frac{k'}{2k} \frac{\text{RIP}_{2k+k'}^-}{\text{RIP}_{k'}^+}} \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_1 - \sqrt{k'} \|\mathcal{A}(\Delta)\| \\ &\geq (1 + \epsilon) \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_1 - \sqrt{k'} \|\mathcal{A}(\Delta)\|. \end{aligned}$$

Putting things together. We are now ready to prove sharpness:

$$\begin{aligned} &\|X^\natural + \Delta\|_1 + \sqrt{k'} \|\mathcal{A}(\Delta)\| - \|X^\natural\|_1 \\ &\geq \left\| \begin{pmatrix} X_{1,1}^\natural & 0 \\ 0 & \Delta_{2,2} \end{pmatrix} \right\|_1 - \|X^\natural\|_1 - \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_1 \\ &\quad + \sqrt{k'} \|\mathcal{A}(\Delta)\| \\ &\geq \|\delta_{k^\perp}\|_1 - \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_1 + \sqrt{k'} \|\mathcal{A}(\Delta)\| \quad (\text{disjoint supports}) \\ &= \left(\frac{2}{2+\epsilon} + \frac{\epsilon}{2+\epsilon} \right) \|\delta_{k^\perp}\|_1 - \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_1 \\ &\quad + \sqrt{k'} \|\mathcal{A}(\Delta)\| \\ &\geq \left(\frac{2}{2+\epsilon} (1+\epsilon) - 1 \right) \left\| \begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} \right\|_1 + \frac{\epsilon}{2+\epsilon} \|\delta_{k^\perp}\|_1 \quad (\text{bound on } \|\delta_{k^\perp}\|_1) \\ &\quad + \left(1 - \frac{2}{2+\epsilon} \right) \sqrt{k'} \|\mathcal{A}(\Delta)\| \\ &\geq \frac{\epsilon}{2+\epsilon} \|\Delta\|_1. \end{aligned}$$

This proves the claim (8) as $\Delta \in \mathbb{R}^{n \times N}$ was arbitrary.

3.2 Sharpness in sparse vector recovery

The proof can be done by following the proof for matrix sensing by treating the signal x^\natural and the sensing vectors in \mathcal{A} as diagonal matrices. The better constant $c = 1$ is obtained by noting there is no extra off-diagonal term for the matrices considered in the argument. Thus, the matrix $\begin{pmatrix} \Delta_{1,1} & \Delta_{1,2} \\ \Delta_{2,1} & 0 \end{pmatrix} = \begin{pmatrix} \Delta_{1,1} & \\ & 0 \end{pmatrix}$ has rank bounded by k instead of $2k$. Alternatively, we give a direct proof in Appendix B.

3.3 Sharpness in covariance estimation

We may repeat the proof of Section 3.1 verbatim after replacing $V = (\mathbb{R}^{n \times N}, \|\cdot\|_1)$ by $V = (\mathbb{S}^n, \|\cdot\|_1)$. We deduce that, under the assumptions of this proposition,

$$\|X\|_1 + \sqrt{k'} \|\mathcal{A}(X) - b\|$$

is $\frac{\epsilon}{2+\epsilon}$ sharp around X^\natural . Next, note that $\text{tr}(X) + 2 \text{dist}(X, \mathbb{S}_+^n) \geq \|X\|_1$. We conclude that

$$\text{tr}(X) + 2 \text{dist}(X, \mathbb{S}_+^n) + \sqrt{k'} \|\mathcal{A}(X) - b\|$$

is $\frac{\epsilon}{2+\epsilon}$ sharp around X^\natural .

3.4 Proof of Lipschitz continuity

We need to show that

$$f(x) + \sqrt{k'} \|\mathcal{A}(x) - b\| + 2 \operatorname{dist}(x, K) \quad (9)$$

is $3 + \sqrt{k'} \operatorname{RIP}_1^+(\mathcal{A})$ Lipschitz with respect to $\|\cdot\|_1$.

By the triangle inequality, the objective functions $f(x) = \|x\|_1$ and $f(X) = \operatorname{tr}(X)$ are both 1-Lipschitz in terms of $\|\cdot\|_1$. In the covariance estimation setting, the distance function $\operatorname{dist}(X, \mathbb{S}_+^n)$ is also 1-Lipschitz in terms of $\|\cdot\|_1$ by the triangle inequality.

Next, we show the function $\|\mathcal{A}(x) - b\|$ is RIP_1^+ Lipschitz. Indeed, for any $x_1, x_2 \in V$, let $\Delta = x_1 - x_2 = \sum_{i=1}^n \Delta_i$, where Δ_i are 1-sparse or has rank no more than 1 (due to singular value decomposition). We have

$$\begin{aligned} & \left| \|\mathcal{A}(x_1) - b\| - \|\mathcal{A}(x_2) - b\| \right| \\ & \leq \|\mathcal{A}(x_1 - x_2)\| = \left\| \mathcal{A} \left(\sum_i^n \Delta_i \right) \right\| \leq \sum_{i=1}^n \|\mathcal{A} \Delta_i\| \\ & \stackrel{(a)}{\leq} \sum_{i=1}^n \operatorname{RIP}_1^+ \|\Delta_i\|_2 \stackrel{(b)}{=} \sum_{i=1}^n \operatorname{RIP}_1^+ \|\Delta_i\|_1 = \operatorname{RIP}_1^+ \|\Delta\|_1. \end{aligned}$$

Here step (a) and (b) are due to the RIP and $\Delta_i, i = 1, \dots, n$ are 1-sparse or have rank no more than 1.

Since $f(x)$ and $\operatorname{dist}(x, K)$ are 1 Lipschitz and $\|\mathcal{A}(x) - b\|$ is RIP_1^+ Lipschitz, our proof for that the function in (9) is $3 + \sqrt{k'} \operatorname{RIP}_1^+(\mathcal{A})$ -Lipschitz is complete.

4 Conditioning, sensing models, and sample complexity

In this section, we describe different sensing models of \mathcal{A} and show that once m exceeds certain thresholds, Problem (P) is well-conditioned in terms of ℓ_1 norm with $W = (\mathbb{R}^m, \|\cdot\|_p)$ where $p = 1, 2$.

To prove this result, we first describe the precise sensing model in Section 4.1. Then we collect bounds on the RIP from the literature that ensure that the premise of Proposition 1 is satisfied. Due to an additional technicality in one of the sensing models (labeled ‘‘Covariance Estimation I’’ below), we give a separate proof of its well-conditioning in Section 4.2. In this section, we equip \mathbb{S}^n and $\mathbb{R}^{n \times N}$ with the trace inner product and equip \mathbb{R}^n with the dot product so that the dual map \mathcal{A}^* can be identified with its adjoint map. A summary of the sensing model, the thresholds, and the norms on $W = \mathbb{R}^m$ is described in Table 1. Our main result of this section is the following theorem.

Theorem 1. *Suppose the space W and the sampling map \mathcal{A} are described according to one of the scenarios in Table 1 and the signal x^\natural is either k -sparse or has rank no more than k . If $m \geq CT(n, N, k)$, where $T(n, N, k)$ is defined in Table 1 and C is a numerical constant, then there are numerical constants $c_1, c_2 > 0$ such that w.h.p. the optimization problem (P) is $(\frac{1}{2}, \sqrt{c_1 k}, 2)$ -sharp around x^\natural and $F_{\sqrt{c_1 k}, 2}$ has a Lipschitz constant bounded by $\sqrt{c_2 k}$. Consequently, (P) has a condition number κ bounded by $2\sqrt{c_2 k}$.*

4.1 Sensing model, sample complexity, and proof via RIP

In this section, we first describe the sensing models and the thresholds $T(n, N, k)$ on m so that bounds on RIP hold for \mathcal{A} when $W = \mathbb{R}^m$ is equipped with either the ℓ_1 or ℓ_2 norms. Then, we

Task	$\mathcal{A}^*(e_i)$	$T(n, N, k)$	norms of $W = \mathbb{R}^m$
Sparse vector recovery	$a_i \in \mathbb{R}^n$	$k \log(n/k)$	ℓ_1 or ℓ_2
Low rank matrix sensing I	$A_i \in \mathbb{R}^{n \times N}$	$\max\{n, N\}k$	ℓ_1 or ℓ_2
Low rank matrix sensing II	$a_i b_i^\top \in \mathbb{R}^{n \times N}$	$\max\{n, N\}k$	ℓ_1
Covariance estimation I	$a_i a_i^\top \in \mathbb{S}^n$	nk	ℓ_1
Covariance estimation II	$a_i a_i^\top - b_i b_i^\top \in \mathbb{S}^n$	nk	ℓ_1

Table 1: Description of different statistical signal recovery tasks and the thresholds for well-conditioning. The entries of A_i , a_i , and b_i are i.i.d. Gaussian random variables with appropriate scaling (see Section 4.1). The conditioning of (P) is measured by $(\frac{1}{2}, \sqrt{c_1 k}, 2, \sqrt{c_2 k})$ where $c_1, c_2 > 0$ are numerical constants.

prove Theorem 1 by verifying the premise of Proposition 1.

Recall that $\text{RIP}_{k'}^-(\mathcal{A})$ and $\text{RIP}_{k'}^+(\mathcal{A})$ are any uniform lower and upper bound on $\|\mathcal{A}x\| / \|x\|_2$ as x ranges over elements of V with support or rank bounded by k' . They will be set to numerical constants c_1, c_2 below and could differ for different sensing models. The norm of $W = \mathbb{R}^m$ will be either the ℓ_1 norm or ℓ_2 norm below.

Sparse vector recovery. For sparse vector recovery, the measurements are of the form

$$\mathcal{A}^*(e_i) = a_i, \quad i = 1, \dots, m,$$

where the measurement vectors $a_i \stackrel{\text{i.i.d.}}{\sim} N(0, I/m)$ in the ℓ_2 setting and $a_i \stackrel{\text{i.i.d.}}{\sim} N(0, I/m^2)$ in the ℓ_1 setting. In both settings, we may set w.h.p. $\text{RIP}_{k'}^+(\mathcal{A}) = c_2$ and $\text{RIP}_{k'}^-(\mathcal{A}) = c_1$ as long as $m \gtrsim k' \log(\frac{n}{k'})$,¹⁰ where c_1 and c_2 are constants independent of k' satisfying $c_2/c_1 \leq 1.1$ and $c_2 \geq 1$. This fact is proved in the ℓ_2 setting following [13, 31] using results on singular values of random matrices; a simple proof can be found in [5, Theorem 5.2]. This fact is proved in the ℓ_1 setting using [47, Lemma 2.1] and [55, Lemma 4.4].

Matrix sensing I. For this scenario of matrix sensing, the measurements are of the form

$$\mathcal{A}^*(e_i) = A_i, \quad i = 1, \dots, m,$$

where each measurement matrix $A_i \in \mathbb{R}^{n \times N}$ has Gaussian entries. Each entry of each matrix is sampled i.i.d. according to $N(0, 1/m)$ in the ℓ_2 setting, and according to $N(0, 1/m^2)$ in the ℓ_1 setting. In both settings, we may set w.h.p. $\text{RIP}_{k'}^+(\mathcal{A}) = c_2$ and $\text{RIP}_{k'}^-(\mathcal{A}) = c_1$ as long as $m \gtrsim k' \max(n, N)$, where c_1 and c_2 are constants independent of k' satisfying $c_2/c_1 \leq 1.1$ and $c_2 \geq 1$. This fact is proved in the ℓ_2 setting in [12, Theorem 2.3]. This fact is proved in the ℓ_1 setting in [41, Proposition 1] and [18, Theorem 6.4].

Matrix sensing II. For this version of matrix sensing, which is more commonly known as bilinear sensing, the measurements are of the form

$$\mathcal{A}^*(e_i) = a_i b_i^\top, \quad i = 1, \dots, m,$$

and the measurement vectors $a_i \stackrel{\text{i.i.d.}}{\sim} N(0, I_n/m)$ and $b_i \stackrel{\text{i.i.d.}}{\sim} N(0, I_N/m)$. We equip $W = \mathbb{R}^m$ with the ℓ_1 norm. In this setting, we may set w.h.p. $\text{RIP}_{k'}^+(\mathcal{A}) = c_2$ and $\text{RIP}_{k'}^-(\mathcal{A}) = c_1$ as long as

¹⁰We write $a \gtrsim b$ if $a \geq Cb$ for some numerical constant $C > 0$.

$m \gtrsim k' \max(n, N)$, where c_1 and c_2 are constants independent of k' satisfying $c_2/c_1 \leq 4$ and $c_2 \geq 1$. This fact is proved according to [9, Theorem 2.2].

Covariance estimation I. For this scenario of covariance estimation, the measurements are of the form

$$\mathcal{A}^*(e_i) = a_i a_i^\top, \quad i = 1, \dots, m,$$

where each measurement vector $a_i \stackrel{\text{i.i.d.}}{\sim} N(0, I_n/m)$. We equip $W = \mathbb{R}^m$ with the ℓ_1 norm. The proof of Theorem 1 for this setting differs from the proof of Theorem 1 for all other settings. This difference stems from the fact that $\langle \mathcal{A}^*(e_i), X^\natural \rangle$ does not have zero mean, thus biasing the output vector as discussed in [21, Section III.B]. Attempting to follow the same proof strategy will need $\text{RIP}_{k'}^+(\mathcal{A})$ to be a numerical constant. However, the quantity $\text{RIP}_{k'}^+(\mathcal{A})$ scales as $\sqrt{k'}$. This prevents us from applying Proposition 1 directly in this setting.

We provide a separate proof for Theorem 1 (Covariance estimation I) in Section 4.2. The proof in this setting is completed by analyzing the conditioning of a related model of covariance estimation (Covariance estimation II).

Covariance estimation II. For this scenario of covariance estimation, the measurements are of the form

$$\mathcal{A}^*(e_i) = a_i a_i^\top - b_i b_i^\top, \quad i = 1, \dots, m,$$

where each measurement vector pair $a_i, b_i \stackrel{\text{i.i.d.}}{\sim} N(0, I_n/(2m))$. We equip $W = \mathbb{R}^m$ with the ℓ_1 norm. In this setting, we may set w.h.p. $\text{RIP}_{k'}^+(\mathcal{A}) = c_2$ and $\text{RIP}_{k'}^-(\mathcal{A}) = c_1$ as long as $m \gtrsim nk'$. Here, c_1 and c_2 are constants independent of k' satisfying $c_2/c_1 \leq 4$ and $c_2 \geq 1$. This fact is proved in [9, Theorem 2.2] and [18, Theorem 6.4].

The Gaussian assumption on the sensing vectors or matrices are assumed for simplicity. One can relax this condition to sub-Gaussian distributions as done in [5, Theorem 5.2], [47, Lemma 2.1], and [18, Theorem 6.4].

Proof of Theorem 1 for all settings except covariance estimation I. As described above, for all settings except covariance estimation I, as long as $m \gtrsim T(n, N, k')$, we may set w.h.p. $\text{RIP}_{k'}^+(\mathcal{A}) = c_2$ and $\text{RIP}_{k'}^-(\mathcal{A}) = c_1$. By setting $k' = \sqrt{C_1 k}$ for a large numerical constant C_1 , we see the premise of Proposition 1 is satisfied and our proof is complete. ■

4.2 Well-conditioning of covariance estimation I

We provide a separate argument for Theorem 1 in the setting of Covariance estimation I.

In this setting, we have that $a_i \sim N(0, I_n/m)$ are i.i.d. We equip the space $W = \mathbb{R}^m$ with the ℓ_1 norm. Our goal is to show that (P) is sharp in terms of ℓ_1 norm with parameters $(\frac{1}{2}, \frac{1}{2}\sqrt{c_1 k}, 2)$ and Lipschitz continuous with Lipschitz constant $L = \sqrt{c_2 k}$ w.h.p. once $m \gtrsim nk$. We will do so by comparing Covariance estimation I with Covariance estimation II.

First, we replace the linear constraint in Problem (P) by $\mathcal{B}(X) = d$ where $\mathcal{B} : \mathbb{S}^n \rightarrow \mathbb{R}^{\lfloor m/2 \rfloor}$, and

$$\mathcal{B}^*(e_i) = \frac{1}{2} a_{2i-1} a_{2i-1}^\top - \frac{1}{2} a_{2i} a_{2i}^\top \quad d_i = \frac{1}{2} b_{2i-1} - \frac{1}{2} b_{2i}, \quad i = 1, \dots, \lfloor m/2 \rfloor. \quad (10)$$

This is distributed as an instance of Covariance estimation II. By Theorem 1 (Covariance estimation II), we know that once $m \gtrsim nk$, it holds that

$$\mathrm{tr}(X) + \sqrt{c_1 k} \|\mathcal{B}(X) - d\|_1 + 2 \mathrm{dist}(X, \mathbb{S}_+^n) - \mathrm{tr}(X^\natural) \geq \frac{1}{2} \|X - X^\natural\|_1. \quad (11)$$

Note that for each $i = 1, \dots, \lfloor m/2 \rfloor$, by the construction (10), we have

$$|\langle \mathcal{B}^*(e_i), X \rangle - d_i| \leq \frac{1}{2} |\langle \mathcal{A}^*(e_{2i-1}), X \rangle - b_{2i-1}| + \frac{1}{2} |\langle \mathcal{A}^*(e_{2i}), X \rangle - b_{2i}|.$$

Combining this fact with (11), we see that the function $\mathrm{tr}(X) + \frac{1}{2} \sqrt{c_1 k} \|\mathcal{A}(X) - b\|_1 + 2 \mathrm{dist}(X, \mathbb{S}_+^n)$ is $\frac{1}{2}$ sharp around X^\natural as well.

To prove the Lipschitz constant is bounded, we utilize [14, Lemma 3.1], which shows that w.h.p.

$$\|\mathcal{A}(X)\| \leq 1.1 \|X\|_1.$$

Hence we see $\mathrm{tr}(X) + \frac{1}{2} \sqrt{c_1 k} \|\mathcal{A}(X) - b\|_1 + 2 \mathrm{dist}(X, \mathbb{S}_+^n)$ is $3 + \sqrt{c_1 k}$ Lipschitz with respect to the ℓ_1 norm. This completes the proof of Theorem 1 (Covariance estimation I).

5 Sharp problem formulations in the presence of noise

In this section, we show that the sharpness of a problem (P) in the noiseless setting $b = \mathcal{A}(x^\natural)$ provides (algorithmically useful) information even in the noisy setting, where $b = \mathcal{A}(x^\natural) + \delta$ with δ small or sparse. We begin with the case where δ is small.

Proposition 2. *Suppose (P) is (μ, r, ℓ) sharp around x^\natural . Let $\delta \in W$ and set $\tilde{b} = b + \delta$.*

- If \tilde{x} minimizes

$$\tilde{F}_{r,\ell}(x) := f(x) + r \|\mathcal{A}(x) - \tilde{b}\| + \ell \mathrm{dist}(x, \mathcal{K}),$$

$$\text{then } \|\tilde{x} - x^\natural\| \leq \frac{2r}{\mu} \|\delta\|.$$

- If \tilde{x} minimizes

$$\tilde{F}_{r,\ell}^{\mathrm{thresh}}(x) := \max\left(\tilde{F}_{r,\ell}(x), F_{r,\ell}(x^\natural) + 3r \|\delta\|\right),$$

$$\text{then } \|\tilde{x} - x^\natural\| \leq \frac{4r}{\mu} \|\delta\|. \text{ Furthermore, } \tilde{F}_{r,\ell}^{\mathrm{thresh}} \text{ is } \frac{\mu}{2}\text{-sharp around its optimizers.}$$

Remark 3. Suppose $F_{r,\ell}$ is L -Lipschitz with $f(x) = \|x\|$, $\mathcal{K} = V$ and $L \geq 1$ as in sparse vector recovery or low-rank matrix sensing. Then we have

$$L \|\|x^\natural\| \geq |F_{r,\ell}(0) - F_{r,\ell}(x^\natural)| = |r \|b\| - \|x^\natural\| \implies (L+1) \|x^\natural\| \geq r \|b\|.$$

Hence, combining this inequality with the first item of Proposition 2, we have

$$\frac{\|\tilde{x} - x^\natural\|}{\|x^\natural\|} \leq \frac{2L+2}{\mu} \frac{\|\delta\|}{\|b\|}.$$

Thus we see that indeed the condition number of (P) controls the relative change of the solution to the relative perturbation to the data vector b :

$$\underbrace{\frac{\|\tilde{x} - x^\natural\|}{\|x^\natural\|}}_{\text{Relative change in solution}} \leq \mathcal{O}(\kappa) \cdot \underbrace{\frac{\|\delta\|}{\|b\|}}_{\text{Relative change in data vector}}. \quad \square$$

Remark 4. When δ is nonzero, the penalization formulation $\tilde{F}_{r,\ell}$ is not necessarily a sharp function. However, the above proposition asserts that $\tilde{F}_{r,\ell}^{\text{thresh}}(x)$ is still sharp. Hence, we may hope to apply the methods described in Section 6, which apply to sharp functions: On the surface, evaluating the function $\tilde{F}_{r,\ell}^{\text{thresh}}(x)$ (and its subgradients) requires access to both x^\natural and $\|\delta\|$. While we will not have access to these quantities in practice, that is of little consequence if we only plan to apply first-order methods (as we suggest in Section 6). Indeed, any first-order method applied to $\tilde{F}_{r,\ell}$ will behave equivalently to the first-order method applied to $\tilde{F}_{r,\ell}^{\text{thresh}}$ until an iterate \tilde{x} satisfying $\|\tilde{x} - x^\natural\| \leq \frac{4r}{\mu} \|\delta\|$ is found. In particular, the algorithms presented in Section 6 applied to $\tilde{F}_{r,\ell}$ will converge linearly to such a point with a rate depending on $\frac{\mu}{2}$. We emphasize that such a procedure is *adaptive* (to the noise level $\|\delta\|$) in both $\|\tilde{x} - x^\natural\|$ and the rate of convergence to \tilde{x} . \square

Proof of Proposition 2. For notational convenience, in this proof we will drop all subscripts r, ℓ .

By the triangle inequality, for all $x \in V$, we have

$$\left| F(x) - \tilde{F}(x) \right| = \left| r \|\mathcal{A}(x) - b\| - r \|\mathcal{A}(x) - \tilde{b}\| \right| \leq r \|\delta\|.$$

For the first claim, note that by optimality of \tilde{x} in \tilde{F} , we have that

$$F(\tilde{x}) \leq \tilde{F}(\tilde{x}) + r \|\delta\| \leq \tilde{F}(x^\natural) + r \|\delta\| \leq F(x^\natural) + 2r \|\delta\|.$$

Combining this inequality with μ -sharpness of F around x^\natural proves the first claim.

Consider the second claim. Note that $\tilde{F}(x^\natural) \leq F(x^\natural) + r \|\delta\|$ so that \tilde{F} achieves values bounded above by the threshold value of $F(x^\natural) + 3r \|\delta\|$. Thus, the set of minimizers of $\tilde{F}^{\text{thresh}}(x)$ is given by $\mathcal{X} := \{x \in V : \tilde{F}(x) \leq F(x^\natural) + 3r \|\delta\|\}$. Then, if \tilde{x} minimizes $\tilde{F}^{\text{thresh}}(x)$, we must have

$$F(\tilde{x}) \leq \tilde{F}(\tilde{x}) + r \|\delta\| \leq F(x^\natural) + 4r \|\delta\|.$$

Combining this inequality with μ -sharpness of F around x^\natural shows that $\|\tilde{x} - x^\natural\| \leq \frac{4r}{\mu} \|\delta\|$.

It remains to show that $\tilde{F}^{\text{thresh}}(x)$ is $\mu/2$ sharp around its optimizers \mathcal{X} . By the definition of sharpness (see Definition 1), the goal is to show that for any $\bar{x} \in V \setminus \mathcal{X}$, there exists $\tilde{x} \in \mathcal{X}$ satisfying

$$\frac{\mu}{2} \|\bar{x} - \tilde{x}\| \leq \tilde{F}^{\text{thresh}}(\bar{x}) - \tilde{F}^{\text{thresh}}(\tilde{x}).$$

Note that for any $\bar{x} \in V \setminus \mathcal{X}$ and $\tilde{x} \in \mathcal{X}$, we have $\tilde{F}^{\text{thresh}}(\bar{x}) = \tilde{F}(\bar{x})$ and $\tilde{F}^{\text{thresh}}(\tilde{x}) = F(x^\natural) + 3r \|\delta\|$.

Set $\tilde{x} = (1 - \alpha)x^\natural + \alpha\bar{x}$ where

$$\alpha = \frac{2r \|\delta\|}{\tilde{F}(\bar{x}) - F(x^\natural) - r \|\delta\|}.$$

As $\tilde{F}(\bar{x}) > F(x^\natural) + 3r \|\delta\|$, we have that α is well-defined and $\alpha \in [0, 1]$. By convexity of \tilde{F} ,

$$\begin{aligned} \tilde{F}(\tilde{x}) &\leq (1 - \alpha)\tilde{F}(x^\natural) + \alpha\tilde{F}(\bar{x}) \\ &\leq (1 - \alpha)(F(x^\natural) + r \|\delta\|) + \alpha\tilde{F}(\bar{x}) \\ &= F(x^\natural) + r \|\delta\| + \left(\frac{2r \|\delta\|}{\tilde{F}(\bar{x}) - F(x^\natural) - r \|\delta\|} \right) (\tilde{F}(\bar{x}) - F(x^\natural) - r \|\delta\|) \\ &= F(x^\natural) + 3r \|\delta\|. \end{aligned}$$

We deduce that $\tilde{x} \in \mathcal{X}$.

Finally, we have

$$\begin{aligned}
\frac{\mu}{2} \|\bar{x} - \tilde{x}\| &= (1 - \alpha) \frac{\mu}{2} \|x^\natural - \bar{x}\| \\
&\leq \frac{1 - \alpha}{2} \left(F(\bar{x}) - F(x^\natural) \right) \\
&\leq \frac{1}{2} \left(\frac{\tilde{F}(\bar{x}) - F(x^\natural) - 3r \|\delta\|}{\tilde{F}(\bar{x}) - F(x^\natural) - r \|\delta\|} \right) \left(\tilde{F}(\bar{x}) - F(x^\natural) + r \|\delta\| \right) \\
&= \frac{1}{2} \left(\tilde{F}(\bar{x}) - F(x^\natural) - 3r \|\delta\| \right) \left(\frac{\tilde{F}(\bar{x}) - F(x^\natural) - 3r \|\delta\| + 4r \|\delta\|}{\tilde{F}(\bar{x}) - F(x^\natural) - 3r \|\delta\| + 2r \|\delta\|} \right) \\
&\leq \tilde{F}(\bar{x}) - F(x^\natural) - 3r \|\delta\|.
\end{aligned}$$

Here, the second line follows by μ -sharpness of F , the third line follows by the definition of α , and the final line follows from the premise that $\tilde{F}(\bar{x}) > F(x^\natural) + 3r \|\delta\|$. \blacksquare

The next proposition shows that if W is equipped with the ℓ_1 norm and **(P)** is sharp in the noiseless case, then exact recovery continues to be possible (with linearly convergent algorithms) in the presence of grossly-but-sparsely-corrupted observations.

Proposition 3. *Suppose **(P)** with $b = \mathcal{A}(x^\natural)$ is (μ, r, ℓ) sharp around x^\natural . Let \mathcal{U} be a finite-dimensional normed real vector space. Let $\mathcal{B} : V \rightarrow \mathcal{U}$ be a linear operator with $r \|\mathcal{B}\| < \mu$ and $\delta \in \mathcal{U}$. Then, the function*

$$F_{r,\ell}(x) + r \|\mathcal{B}(x) - \mathcal{B}(x^\natural) - \delta\| \quad (12)$$

is $(\mu - r \|\mathcal{B}\|)$ -sharp around x^\natural .

Proof. Let $\bar{\delta} \in \mathcal{U}^*$ be a vector such that $\|\bar{\delta}\| = 1$ and $\langle \bar{\delta}, \delta \rangle = \|\delta\|$. Then, $\partial \|\mathcal{B}(x) - \mathcal{B}(x^\natural) - \delta\|$ evaluated at x^\natural contains $\mathcal{B}^*(\bar{\delta})$. By additivity of the subgradient, the subgradient of (12) evaluated at x^\natural contains $\partial F_{r,\ell}(x^\natural) + r \mathcal{B}^*(\bar{\delta})$. This set in turn contains $B_{V^*}(0, \mu - r \|\mathcal{B}\|)$ by the triangle inequality, Lemma 1, and the fact that $\|\mathcal{B}^*\| = \|\mathcal{B}\|$. We conclude that $F_{r,\ell}(x) + r \|\mathcal{B}(x) - \mathcal{B}(x^\natural) - \delta\|$ is $(\mu - r \|\mathcal{B}\|)$ sharp around x^\natural . \blacksquare

Example 1. Consider the phase retrieval problem where an α -fraction of the observations are corrupted arbitrarily. Formally, consider the following procedure: Let $V = (\mathbb{S}^n, \|\cdot\|_*)$ and $\tilde{W} = (\mathbb{R}^{\tilde{m}}, \|\cdot\|_1)$. Let $m = \lceil (1 - \alpha)\tilde{m} \rceil$. Fix $X^\natural \in \mathbb{S}_+^n \subseteq V$ to be rank-one and let $\tilde{\mathcal{A}} : V \rightarrow \tilde{W}$ be the random linear map

$$\tilde{\mathcal{A}}(X)_i = g_i^\top X g_i, \quad \text{where } g_i \sim N(0, I/\tilde{m}).$$

Let $\delta \in \tilde{W}$ denote an arbitrary vector, chosen possibly adversarially, with only the guarantee that $\text{supp}(\delta) \subseteq [m + 1, \tilde{m}]$. Set $\tilde{b} = \tilde{\mathcal{A}}(X^\natural) + \delta$. Our goal is to recover X^\natural from $\tilde{\mathcal{A}}$ and \tilde{b} .

Let W, U denote the decomposition of \tilde{W} along the coordinates $[m]$ and $[m + 1, \tilde{m}]$. Let $\mathcal{A} : V \rightarrow W$ and $\mathcal{B} : V \rightarrow U$ denote the corresponding restrictions of $\tilde{\mathcal{A}}$. Slightly abusing notation, we will let $\delta \in U$ denote the restriction of $\delta \in \tilde{W}$.

Now, suppose that

$$\min_{X \in \mathbb{S}_+^n} \left\{ \text{tr}(X) : \begin{array}{l} \mathcal{A}(X) = \mathcal{A}(X^\natural) \\ X \in \mathbb{S}_+^n \end{array} \right\}$$

is (μ, r, ℓ) sharp around X^\natural . As \tilde{W} carries the ℓ_1 norm and W, U are coordinate subspaces, we may write

$$\begin{aligned} & \text{tr}(X) + r \left\| \tilde{\mathcal{A}}(X) - \tilde{b} \right\| + \ell \text{dist}(X, \mathbb{S}_+^n) \\ &= \text{tr}(X) + r \left\| \mathcal{A}(X) - \mathcal{A}(X^\natural) \right\| + r \left\| \mathcal{B}(X) - \mathcal{B}(x^\natural) - \delta \right\| + \ell \text{dist}(X, \mathbb{S}_+^n). \end{aligned} \tag{13}$$

Proposition 3 states that if $\mu > r \|\mathcal{B}\|$, then X^\natural is the unique minimizer of (13) despite the corruption δ . By [57], we know that w.h.p., $\|\mathcal{B}\| \leq \left(\frac{\sqrt{\alpha\tilde{m}} + 2\sqrt{n}}{\sqrt{\tilde{m}}} \right)^2 \leq 2\alpha + 8(n/\tilde{m})$. Combining these bounds, we have that X^\natural is the unique minimizer of the sharp unconstrained minimization problem (13) if $\alpha \lesssim \mu/r$ and $n/\tilde{m} \lesssim \mu/r$. In particular, μ/r controls the fraction of allowed gross corruption in the observation $\tilde{\mathcal{A}}(X^\natural)$. \square

6 First-order methods for non-Euclidean sharp minimization

A natural class of algorithms for large-scale problems of the form (P) with sharpness is the restarted mirror descent (RMD) algorithm (Algorithm 2). This algorithm generalizes similar algorithms for minimizing sharp functions in a Euclidean norm [35, 48, 59] and has nearly dimension-independent linear convergence rates that depend explicitly on sharpness (see Theorem 2). Algorithm 2 can be applied to $F_{r,\ell}$, the sharp exact penalty formulation of (P), or any of its sharp perturbations (see Propositions 2 and 3).

We describe and analyze the basic RMD scheme in Section 6.1, and discuss two variants of RMD in Section 6.2.

6.1 Restarted Mirror Descent

We restrict our attention to sharp Lipschitz convex functions in an ℓ_p or Schatten p -norm for $p \in [1, 2]$: Throughout this section, let V be a normed finite-dimensional real vector space. We will overload notation so that we can simultaneously consider three separate settings:

Assumption 1. Let $p \in [1, 2]$ and either:

- Let $V = (\mathbb{R}^n, \|\cdot\|_p)$ where $\|\cdot\|_p$ is the ℓ_p norm. Or,
- let $V = (\mathbb{S}^n, \|\cdot\|_p)$ where $\|\cdot\|_p$ is the Schatten p -norm. Or,
- let $V = (\mathbb{R}^{n \times N}, \|\cdot\|_p)$ where $\|\cdot\|_p$ is the Schatten p -norm.

If we are in the third case, we will assume that $n \leq N$. \square

Note that the Schatten 1-norm, $\|\cdot\|_1$ coincides with the nuclear norm.

Recall the mirror descent algorithm and its guarantee [8, Theorem 4.2].

Algorithm 1 Mirror Descent

Given $f : V \rightarrow \mathbb{R}$, $\bar{x} \in \mathbb{R}^d$, $\eta > 0$, $T \in \mathbb{N}$, $h : V \rightarrow \mathbb{R}$

- Let $x_0 = \bar{x}$, $\theta_0 = 0 \in \mathbb{R}^d$
 - For $t = 1, \dots, T$
 - Let $g_t \in \partial f(x_{t-1})$ and set $\theta_t = \theta_{t-1} - \eta \cdot g_t$
 - Set $x_t = (\nabla h)^{-1}(\theta_t)$
 - Output the x_t minimizing $f(x_t)$ among $t \in [0, T]$.
-

Lemma 3. *Let $h : V \rightarrow \mathbb{R}$ be differentiable and σ -strongly convex with respect to the norm on V . Suppose $f : V \rightarrow \mathbb{R}$ is convex and L -Lipschitz with respect to the norm on V . Then, the mirror descent algorithm initialized at \bar{x} , run for t iterations, with step-size η , produces \tilde{x} such that*

$$f(\tilde{x}) - f(x^*) \leq \frac{L^2\eta}{2\sigma} + \frac{D_h(x^*|\bar{x})}{\eta t}$$

where $x^* \in V$ is any minimizer of f . Here, each iteration requires computing a single subgradient of f and applying $(\nabla h)^{-1}$ and arithmetic operations in V and V^* .

In the bound above, the quantity $D_h(\cdot|\cdot)$ is a Bregman divergence term. We elaborate on this term for our choice of the map h . Let $p \in (1, 2]$ and define $h_{\bar{x}}(x) := \frac{1}{2} \|x - \bar{x}\|_p^2$ for $\bar{x} \in V$. It is known [4] that in each of the three setups in Assumption 1, that $h_{\bar{x}}$ is differentiable and $(p-1)$ -strongly convex w.r.t. $\|\cdot\|_p$. Furthermore, for all $x^* \in V$, the Bregman divergence (associated to $h_{\bar{x}}$) of x^* with respect to \bar{x} is

$$D_{h_{\bar{x}}}(x^*|\bar{x}) := h_{\bar{x}}(x^*) - h_{\bar{x}}(\bar{x}) - \langle \nabla h_{\bar{x}}(\bar{x}), x^* - \bar{x} \rangle = h_{\bar{x}}(x^*) = \frac{1}{2} \|x^* - \bar{x}\|_p^2.$$

Thus, if f is L -Lipschitz w.r.t. $\|\cdot\|_p$, then the output $\tilde{x} = \text{mirror}(h_{\bar{x}}, f, L, \bar{x}, t, \eta)$ has suboptimality bounded by

$$f(\tilde{x}) - f(x^*) \leq \frac{L^2\eta}{2(p-1)} + \frac{\|x^* - \bar{x}\|_p^2}{2\eta T}.$$

This bound holds simultaneously for all minimizers $x^* \in V$ of f .

Remark 5. Consider a setup from Assumption 1 and let $p \in (1, 2]$. In each iteration of mirror descent (applied with $h_{\bar{x}}$), we must evaluate $(\nabla h_{\bar{x}})^{-1}$ on some input $\theta \in V^*$. By [52, Corollary 23.5.1], this is equivalent to evaluating $\nabla \bar{h}_{\bar{x}}(\theta)$ where $\bar{h}_{\bar{x}}$ is the convex conjugate of $h_{\bar{x}}$. Thus,

$$(\nabla h_{\bar{x}})^{-1}(\theta) = \nabla \bar{h}_{\bar{x}}(\theta) = \bar{x} + \nabla \frac{1}{2} \|\theta\|_q^2 = \bar{x} + \frac{\text{sign}(\theta) \circ |\theta|^{q-1}}{\|\theta\|_q^{q-2}}.$$

Here, q is the Hölder dual to p and the expression $\text{sign}(\theta) \circ |\theta|^{q-1}$ is applied entrywise to the entries of θ if θ is a vector and to the singular values of θ if θ is a matrix. \square

Remark 6. In the matrix setting, computing the mirror map requires performing an SVD in each iteration. We expect that this full SVD may be replaced by a partial SVD near an optimal low-rank solution (following [34]) and leave this extension for future work. \square

Now consider the following restarted variant of mirror descent where the step size *and* mirror map h update at each restart.

Algorithm 2 $\text{RMD}(f, L, x_0, K, \{\eta_k\}, t, p)$

- For $k = 1, 2, \dots, K$
 - Set $x_k \leftarrow \text{mirror}(h_{x_{k-1}}, f, L, x_{k-1}, t, \eta_k)$ where

$$h_{x_{k-1}}(x) := \frac{1}{2} \|x - x_{k-1}\|_p^2$$

- Output x_K
-

Theorem 2. Consider a setup from Assumption 1 and let $p \in (1, 2]$. Suppose $f : V \rightarrow \mathbb{R}$ is L -Lipschitz and μ -sharp w.r.t. $\|\cdot\|_p$ and $x_0 \in V$ satisfies $f(x_0) - f^* \leq \epsilon_0$. Let $\epsilon_k = \epsilon_0 e^{-k/2}$, $K = \lceil 2 \ln(\frac{\epsilon_0}{\epsilon}) \rceil$, $t = \lceil \frac{eL^2}{\mu^2(p-1)} \rceil$, and $\eta_k = \frac{(p-1)\epsilon_k}{L^2}$. Then, each iterate x_k in $\text{RMD}(f, L, x_0, K, \{\eta_k\}, t, p)$ satisfies $f(x_k) - f^* \leq \epsilon_k$. In particular, an $0 < \epsilon \leq \epsilon_0$ -optimizer can be computed in

$$O\left(\frac{L^2}{\mu^2(p-1)} \log\left(\frac{\epsilon_0}{\epsilon}\right)\right)$$

total mirror descent steps. If $f : V \rightarrow \mathbb{R}$ is L -Lipschitz and μ -sharp w.r.t. $\|\cdot\|_1$, we may apply the above statement with $p = 1 + \frac{1}{\ln n}$, sharpness μ , and Lipschitz constant eL w.r.t. $\|\cdot\|_p$.

Proof. The claim holds for $k = 0$. Now let $k \geq 1$. By μ -sharpness of f , there exists an optimizer x^* of f , satisfying $\mu \|x_{k-1} - x^*\|_p \leq f(x_{k-1}) - f^*$. By Lemma 3, we have

$$\begin{aligned} f(x_k) - f^* &\leq \frac{L^2 \eta_k}{2(p-1)} + \frac{\|x^* - x_{k-1}\|_p^2}{2\eta_k t} \\ &\leq \frac{L^2 \eta_k}{2(p-1)} + \frac{e\epsilon_k^2}{2\mu^2 \eta_k t} \\ &\leq \frac{L^2 \eta_k}{2(p-1)} + \frac{\epsilon_k^2 (p-1)}{2\eta_k L^2} \\ &= \epsilon_k. \end{aligned}$$

The setting of L -Lipschitz, μ -sharp convex functions w.r.t. $\|\cdot\|_1$ reduces to the setting of $p = 1 + \frac{1}{\ln(n)}$ by the bounds

$$\frac{1}{e} \|w\|_1 \leq \|w\|_p \leq \|w\|_1, \quad \forall w \in V,$$

which hold [4] in each of the setups in Assumption 1. Specifically, these inequalities imply that f is μ -sharp and eL -Lipschitz w.r.t. $\|\cdot\|_p$. ■

Remark 7. In [54], the authors consider restarted versions of various accelerated first order methods for sharp problems including problems in non-Euclidean spaces. In the non-Euclidean setting, [54] suggests algorithms for functions f that satisfy the following proxy for sharpness:

$$f(x) - f(x^\natural) \geq \mu \sqrt{D_h(x^\natural \| x)}. \quad (14)$$

Here, $h : V \rightarrow \mathbb{R}$ is a *fixed* differentiable and σ -strongly convex function¹¹ and D_h is the Bregman divergence induced by h . Then, calculations almost identical to the proof of Theorem 2 show that

¹¹The notation of [54] uses a 1-strongly convex function, but this only changes the value of μ by a factor of $\sqrt{\sigma}$.

restarted mirror descent with the mirror map h in every restart produces an ϵ -suboptimal solution in $O\left(\frac{L^2}{\mu^2\sigma} \log\left(\frac{\epsilon_0}{\epsilon}\right)\right)$ mirror descent steps.

Unfortunately, (14) is *not* implied by sharpness in the original norm and may be overly restrictive, as illustrated in the example below. Specifically, in standard prox setups for mirror descent in ℓ_p norms or Schatten- p norms, (14) cannot hold for any $\mu > 0$ unless x^\natural has full support in the vector case or full rank in the matrix case. \square

Example 2. Let $V = (\mathbb{R}^n, \|\cdot\|_p)$ where $p \in (1, 2)$ and $n \geq 2$. Let $f(x) := \|x - x^\natural\|_p$ where $x^\natural = e_1$. Clearly, f is 1-sharp around x^\natural and 1-Lipschitz w.r.t. $\|\cdot\|_p$.

Two standard prox setups for mirror descent [45] in this setting are to take $h(x) = \frac{1}{2} \|x\|_p^2$ (in the unbounded case) or $h(x) = \frac{1}{p} \|x\|_p^p$ (in the bounded case); see [45, Theorem 2.1] for the corresponding strong convexity parameters. Consider the first setting, i.e., $h(x) = \frac{1}{2} \|x\|_p^2$. Letting $x_\epsilon = x^\natural + \epsilon e_2$, we have by Bernoulli's inequality that

$$\begin{aligned} \sqrt{D_h(x^\natural|x)} &= \sqrt{\frac{1}{2}(1 + |\epsilon|^p)^{2/p} - \frac{1}{2} - \langle e_1, \epsilon e_2 \rangle} \\ &= 2^{-1/2} \sqrt{(1 + |\epsilon|^p)^{2/p} - 1} \\ &\geq p^{-1/2} |\epsilon|^{p/2}. \end{aligned}$$

On the other hand, $f(x_\epsilon) - f(x^\natural) \leq |\epsilon|$. Thus, letting $|\epsilon| \rightarrow 0$, we see that (14) cannot hold for *any* $\mu > 0$.

Next, suppose $h(x) = \frac{1}{p} \|x\|_p^p$. Then,

$$\sqrt{D_h(x^\natural|x)} = \sqrt{\frac{1}{p}(1 + |\epsilon|^p) - \frac{1}{p} - \langle e_1, \epsilon e_2 \rangle} = |\epsilon|^{p/2}.$$

Again, comparing this bound to the fact that $f(x_\epsilon) - f(x^\natural) \leq |\epsilon|$, we deduce that (14) cannot hold for *any* $\mu > 0$. \square

6.2 Variants of RMD

The base version of Algorithm 2 requires knowing the sharpness μ of the function f . We describe two variants below. The first variant, Polyak-RMD (Algorithm 3), assumes instead that the optimal value of f is known and achieves the same convergence rate as Algorithm 2 with the optimal choice of μ . This is the algorithm that we will implement for experimental purposes in Section 8. The second variant, Adaptive-RMD (Algorithm 4), does not require the true optimal value or the sharpness constant, but its convergence rate incurs an additional logarithmic term.

Algorithm 3 Polyak-RMD($f, L, x_0, K, \{\eta_k\}, f^*, \{\epsilon_k\}, p$)

- For $k = 1, 2, \dots, K$
 - Run mirror descent $\text{mirror}(h_{x_{k-1}}, f, L, x_{k-1}, \infty, \eta_k)$ until it finds an iterate x_k satisfying

$$f(x_k) - f^* \leq \epsilon_k$$

- Output x_K
-

The convergence guarantee for Polyak-RMD (Algorithm 3) and its proof are entirely identical to that of Algorithm 2.

Proposition 4. Consider a setup from Assumption 1 and let $p \in (1, 2]$. Suppose $f : V \rightarrow \mathbb{R}$ is L -Lipschitz and μ -sharp w.r.t. $\|\cdot\|_p$ and $x_0 \in V$ satisfies $f(x_0) - f^* \leq \epsilon_0$. Let $\epsilon_k = \epsilon_0 e^{-k/2}$, $K = \lceil 2 \ln(\frac{\epsilon_0}{\epsilon}) \rceil$, and $\eta_k = \frac{(p-1)\epsilon_k}{L^2}$. Then, each iterate x_k in $\text{Polyak-RMD}(f, L, x_0, K, \{\eta_k\}, f^*, \{\epsilon_k\}, p)$ is computed in at most $\lceil \frac{eL^2}{\mu^2(p-1)} \rceil$ mirror descent steps. In particular, an $0 < \epsilon \leq \epsilon_0$ -optimizer can be computed in

$$O\left(\frac{L^2}{\mu^2(p-1)} \log\left(\frac{\epsilon_0}{\epsilon}\right)\right)$$

total mirror descent steps. If $f : V \rightarrow \mathbb{R}$ is L -Lipschitz and μ -sharp w.r.t. $\|\cdot\|_1$, we may apply the above statement with $p = 1 + \frac{1}{\ln n}$, sharpness μ , and Lipschitz constant eL w.r.t. $\|\cdot\|_p$.

The next variant of Algorithm 2 does not require either μ or f^* but incurs an additional logarithmic dependence on ϵ^{-1} . Its structure and proof are inspired by [51].

Algorithm 4 Adaptive-RMD($f, L, \bar{x}, \epsilon, \epsilon_0, p$)

- Set $K = 1 + \lceil \lg(\frac{\epsilon_0}{\epsilon}) \rceil$
- Synchronously and in parallel, run K copies of mirror descent. We will refer to these processes as $\text{mirror}_1, \dots, \text{mirror}_K$. For each $i \in [K]$
 - Initialize $\epsilon_i := \epsilon_0 2^{-i}$, $\eta_i := \frac{(p-1)\epsilon_i}{L^2}$, and $x_i \leftarrow \bar{x}$.
 - Repeat:
 - * Run $\text{mirror}(h_{x_i}, f, L, x_i, \infty, \eta_i)$ until either (i) mirror_i finds a \tilde{x} satisfying

$$f(\tilde{x}) \leq f(x_i) - \epsilon_i$$

or (ii) mirror_i receives \tilde{x} from mirror_{i-1} satisfying $f(\tilde{x}) \leq f(x_i) - \epsilon_i$

- * In either case, update $x_i \leftarrow \tilde{x}$. If $i < K$, also pass x_i to mirror_{i+1} .
-

Proposition 5. Consider a setup from Assumption 1 and let $p \in (1, 2]$. Suppose $f : V \rightarrow \mathbb{R}$ is L -Lipschitz and μ -sharp w.r.t. $\|\cdot\|_p$ and $\bar{x} \in V$ satisfies $f(\bar{x}) - f^* \leq \epsilon_0$. Let $0 < \epsilon < \epsilon_0$. Then, the quantity x_K in Algorithm 4 satisfies $f(x_K) - f^* \leq \epsilon$ after at most $O\left(\frac{L^2}{\mu^2(p-1)} \log\left(\frac{\epsilon_0}{\epsilon}\right)\right)$ rounds of synchronous computation or

$$O\left(\frac{L^2}{\mu^2(p-1)} \log\left(\frac{\epsilon_0}{\epsilon}\right)^2\right)$$

total mirror descent steps. If $f : V \rightarrow \mathbb{R}$ is L -Lipschitz and μ -sharp w.r.t. $\|\cdot\|_1$, we may apply the above statement with $p = 1 + \frac{1}{\ln n}$, sharpness μ , and Lipschitz constant eL w.r.t. $\|\cdot\|_p$.

Proof. For the analysis, we will imagine marking mirror_i as soon as $f(x_i) - f^*$ is small enough: Each process $\text{mirror}_1, \dots, \text{mirror}_K$ begins unmarked. Upon a restart of mirror_i , we will mark mirror_i if

$$f(x_i) - f^* \leq 2\epsilon_i.$$

For concreteness, assume the marking happens after x_i is communicated to mirror_{i+1} (if $i < K$).

Note that once a process is marked, it cannot become unmarked. Furthermore, if all K processes are marked, then $f(x_K) - f^* \leq 2^{-K+1}\epsilon_0 \leq \epsilon$. Our goal is to bound the amount of total work before all K processes become marked.

Note that

$$f(\tilde{x}) \leq f^* + \epsilon_0 = f^* + 2\epsilon_1$$

so that `mirror`₁ is marked at the very beginning of the algorithm.

Suppose after some amount of computation that `mirror`₁, ..., `mirror` _{$i-1$} are all marked and `mirror` _{i} is the first unmarked process for $i > 1$. We claim that $f(x_i) - f^* = O(\epsilon_i)$. The fact that `mirror` _{$i-1$} is marked means that some x_{i-1} satisfying $f(x_{i-1}) - f^* \leq 2\epsilon_{i-1}$ was communicated to `mirror` _{i} in prior computation. Thus, it holds that

$$f(x_i) \leq f(x_{i-1}) + \epsilon_i \leq f^* + \epsilon_i + 2\epsilon_{i-1} = f^* + 5\epsilon_i.$$

While `mirror` _{i} remains unmarked, i.e., while

$$f(x_i) - f^* \geq 2\epsilon_i,$$

we may apply Lemma 3 to bound the number of iterations required for `mirror` _{i} to trigger a restart. Specifically, suppose \tilde{x} is the best iterate seen after T iterations of `mirror` _{i} since its last restart. Then by Lemma 3,

$$\begin{aligned} f(\tilde{x}) - f(x_i) &= (f(\tilde{x}) - f^*) - (f(x_i) - f^*) \\ &\leq \frac{L^2\eta_i}{2(p-1)} + \frac{\|x^* - x_i\|_p^2}{2\eta_i T} - 2\epsilon_i \\ &= \frac{\epsilon_i}{2} + \frac{L^2(f(x_i) - f^*)^2}{2\mu^2(p-1)\epsilon_i T} - 2\epsilon_i \\ &\leq \frac{\epsilon_i}{2} + \frac{L^2(5\epsilon_i)^2}{2\mu^2(p-1)\epsilon_i T} - 2\epsilon_i \\ &= \left(-\frac{3}{2} + \frac{25L^2}{2\mu^2(p-1)T} \right) \epsilon_i. \end{aligned}$$

Thus, after $T = \frac{25L^2}{(p-1)\mu^2}$ iterations within `mirror` _{i} , we have that $f(\tilde{x}) \leq f(x_i) - \epsilon_i$. Of course, a restart of `mirror` _{i} may also be triggered by `mirror` _{$i-1$} . In either case, a restart occurs at least every T iterations of `mirror` _{i} so long as it remains the first unmarked process. Clearly, this can only happen up to three times before `mirror` _{i} itself must become marked.

We conclude that after $3KT$ synchronous iterations or $3K^2T$ total iterations of mirror descent, all K processes are marked. ■

7 Relating sharpness to strict complementarity

In this section, we relate sharpness to a notion of conditioning known as strict complementarity. This notion has been studied extensively for linear-conic optimization problems and linearly constrained convex optimization problems (often in the context of deriving efficient algorithms) [1, 24, 25, 27, 34, 36, 58]. We define a notion of strict complementarity for (P).

Definition 7. We say that *strict complementarity* holds for an instance of (P) if there exists an optimal solution x^\natural for which

$$0 \in \text{rint} \left(\partial f(x^\natural) - \text{range}(\mathcal{A}^*) - \left(\mathcal{K}^* \cap (x^\natural)^\perp \right) \right). \quad \square$$

[30, Definition 4] and [46, Defintion 2] define strict complementarity for a linear–conic optimization problem with facially exposed cones, i.e., (P) where $f(x) = \langle c, x \rangle$ for some $c \in V^*$ and where \mathcal{K} and its dual cone are both facially exposed, as the existence of $y^\natural \in W^*$ such that

$$c - \mathcal{A}^*(y^\natural) \in \text{rint}(\mathcal{K}^* \cap (x^\natural)^\perp). \quad (15)$$

We will treat (15) as a definition of strict complementarity for linear–conic optimization problems even when K or its dual cone are not facially exposed.

[27, Section 4] defines strict complementarity for a linearly constrained convex optimization problem, i.e., (P) where $\mathcal{K} = V$, as the existence of $y^\natural \in W^*$ such that

$$\mathcal{A}^*(y^\natural) \in \text{rint} \left\{ \partial f(x^\natural) \right\}. \quad (16)$$

Lemma 4. *Suppose (P) is (μ, r, ℓ) -sharp around x^\natural . Then, (P) satisfies strict complementarity. Furthermore,*

- *if $f(x) = \langle c, x \rangle$ for some $c \in V^*$, then strict complementarity holds in the sense of (15).*
- *if $\mathcal{K} = V$, then strict complementarity holds in the sense of (16).*

Proof. From Lemma 1, we know that Problem (P) is (μ, r, ℓ) -sharp around x^\natural implies that

$$\begin{aligned} 0 &\in \text{int} \left(\partial f(x^\natural) - \mathcal{A}^*(B_{W^*}(0, r)) - \left(\mathcal{K}^* \cap (x^\natural)^\perp \cap B_{V^*}(0, L) \right) \right) \\ &\subseteq \text{int} \left(\partial f(x^\natural) - \text{range}(\mathcal{A}^*) - \left(\mathcal{K}^* \cap (x^\natural)^\perp \right) \right) \\ &\stackrel{(a)}{=} \text{rint}(\partial f(x^\natural)) - \text{range}(\mathcal{A}^*) - \text{rint} \left(\mathcal{K}^* \cap (x^\natural)^\perp \right). \end{aligned}$$

The step (a) follows as the interior and relative interior for a full-dimensional set are the same, and by the additivity of relative interior. This proves the first claim. Now, further rearranging, we deduce there exists $y^\natural \in W^*$ such that

$$\mathcal{A}^*(y^\natural) \in \text{rint}(\partial f(x^\natural)) - \text{rint}(\mathcal{K}^* \cap (x^\natural)^\perp).$$

The two additional claims follow: In the first case $\partial f(x^\natural) = \{c\}$, and in the second case $\mathcal{K}^* \cap (x^\natural)^\perp = \{0\}$. ■

Combining Lemma 4 with Theorem 1, we deduce that strict complementarity holds in each of the settings described in Section 4. Strict complementarity in the setting of phase retrieval may be surprising as Problem (2), also known as the PhaseLift, is highly degenerate. For example, it is known that w.h.p., the unique feasible solution to (2) will be x^\natural [10] so that Slater’s condition does not hold. In fact, simply understanding when a dual optimal solution exists may require substantial effort [37]. In comparison, Lemma 4 and Proposition 1 together imply that the much stronger condition of sharpness holds under an appropriate RIP condition.

The following two examples illustrate caveats in working with strict complementarity and sharpness. First, strict complementarity does not in general imply sharpness. Second, sharpness depends on the presentation of a problem and may be lost even under standard reformulations that are often considered for nuclear norm minimization.

Example 3. Consider the following nuclear norm minimization problem

$$\min_{X \in \mathbb{R}^{2 \times 2}} \{\|X\|_1 : \mathcal{A}(X) := X_{1,1} = 1\}. \quad (17)$$

This problem has a unique minimizer $X^\natural = E_{1,1}$. We compute for (17):

$$\begin{aligned} \partial f(X^\natural) &= \left\{ \begin{pmatrix} 1 & \\ & \alpha \end{pmatrix} : \alpha \in [-1, 1] \right\}, \quad \text{and} \\ \text{range}(\mathcal{A}^*) &= \left\{ \begin{pmatrix} \beta & \\ & 0 \end{pmatrix} : \beta \in \mathbb{R} \right\}. \end{aligned}$$

Thus, strict complementarity holds in the sense of (16). On the other hand, both sets above are contained in a two-dimensional subspace of $\mathbb{R}^{2 \times 2}$. By Lemma 2, we deduce that (17) is *not* (μ, r, ℓ) -sharp around X^\natural for any $\mu > 0$, $r, \ell \geq 0$. \square

Example 4. Consider a nuclear norm minimization problem

$$\min_{X \in \mathbb{R}^{n \times n}} \{\|X\|_1 : \mathcal{A}(X) = b\} \quad (18)$$

with a unique minimizer $X^\natural \neq 0$. For simplicity, we assume that X^\natural is a nonnegative diagonal square matrix; however, similar calculations hold for any nonzero (possibly non-square) matrix.

The standard SDP reformulation of the operator norm minimization problem is given by

$$\min_{\Xi \in \mathbb{S}^{2n}} \left\{ \text{tr}(\Xi) : \begin{array}{l} \tilde{\mathcal{A}}(\Xi) = b \\ \Xi \in \mathbb{S}_+^{2n} \end{array} \right\}, \quad (19)$$

where $\tilde{\mathcal{A}}(\Xi) := \mathcal{A}(\text{the top-right } n \times n \text{ block of } \Xi)$.

It is easy to verify that (19) has a unique minimizer $\Xi^\natural = \begin{pmatrix} 1 & \\ & 1 \end{pmatrix} \otimes X^\natural$ so that solving (19) allows us to recover X^\natural . Unfortunately, we claim that (19) is *not* sharp: Fix $e \in \text{range}(X^\natural)$ nonzero and note that $\ker(\Xi^\natural) = \mathbb{R}^2 \otimes \ker(X^\natural) + \begin{pmatrix} 1 & \\ & -1 \end{pmatrix} \otimes \mathbb{R}^n$. Then,

$$\begin{aligned} \partial f(\Xi^\natural) &= I, \\ \text{range}(\tilde{\mathcal{A}}^*) &= \left\{ \begin{pmatrix} 0 & \xi \\ \xi^\top & 0 \end{pmatrix} : \xi \in \text{range}(\mathcal{A}^*) \right\}, \\ K^* \cap (\Xi^\natural)^\perp &\subseteq \mathbb{S}^{\ker(\Xi^\natural)} = \text{span} \left\{ zz^\top : z \in \ker(\Xi^\natural) \right\}. \end{aligned}$$

Each of the three sets above is contained in the proper subspace $(\begin{pmatrix} 1 & \\ & -1 \end{pmatrix} \otimes ee^\top)^\perp$. By Lemma 2, we deduce that (19) is not (μ, r, ℓ) -sharp around Ξ^\natural for any $\mu > 0$, $r, \ell \geq 0$.

To summarize, the SDP reformulation (19) is never sharp (even if (18) is). \square

8 Numerical experiments

We corroborate our theoretical results with preliminary experiments on the sparse recovery, low-rank matrix sensing, and phase retrieval problems. These experiments suggest that first-order methods based on restarted mirror descent can converge linearly on a variety of statistical recovery problems. Furthermore, the constants in the convergence rates improve (possibly drastically) with the number of samples and are nearly dimension independent.

All experiments were performed on a 2021 Macbook Pro with an Apple M1 pro chip and 32GB of memory. Our code is publicly available at

https://github.com/alexlihengwang/sharpness_well_conditioning

8.1 Implementation details

We implement Polyak-RMD (Algorithm 3). While the assumption that f^* is known is not reflective of real applications, we believe these experiments highlight *more* clearly the connection between sharpness, sample complexity, and dimension. Indeed, the performance of Polyak-RMD is dependent on the *true* value of the sharpness μ (for fixed r, ℓ). On the other hand, the performance of Algorithm 2 is affected by our provided lower bound of μ and thus depends, in part, on our ability to bound μ .

8.1.1 Sparse recovery

We generate a random vector $x^\natural \in \mathbb{R}^n$ with k nonzero entries and a vector of observations $b = Ax^\natural$ where $A \in \mathbb{R}^{m \times n}$. For this setting, $T = (2k \log(n/k) + 1.25k + 1)$ is the current best estimate of the statistical threshold for sparse recovery [17]. The positions of the nonzero entries in x^\natural are chosen uniformly at random, and the values of these entries are set to independent Gaussian random variables. These entries are further normalized so that $\|x^\natural\|_1 = 1$. The matrix A is set to have independent Gaussian entries with variance $1/m$. We run Polyak-RMD on the function

$$\|x\|_1 + 3\sqrt{k} \|Ax - b\|_2.$$

8.1.2 Low-rank matrix sensing

We generate a random rank- k matrix $X^\natural \in \mathbb{R}^{n \times n}$ and a vector of observations $b = \mathcal{A}(X^\natural)$ where $\mathcal{A} : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^m$. For this setting, $T = 6nk - 3k^2 + 1$ is the current best estimate of the statistical threshold for low-rank matrix sensing [17]. The matrix X^\natural is set to a rank- k truncation of a random $n \times n$ matrix with independent Gaussian entries. We normalize $\|X^\natural\|_1 = 1$. The linear operator \mathcal{A} is defined as $\mathcal{A}(X) = (\langle G_i, X \rangle)_i$ where $G_i \in \mathbb{R}^{n \times n}$ has independent Gaussian entries for each $i \in [m]$ with variance $1/m$. We run Polyak-RMD on the function

$$\|X\|_1 + 3\sqrt{k} \|\mathcal{A}(X) - b\|_2.$$

8.1.3 Phase retrieval

We generate a random rank-one PSD matrix $X^\natural \in \mathbb{S}^n$ and a vector of observations $b = \mathcal{A}(X^\natural)$ where $\mathcal{A} : \mathbb{S}^n \rightarrow \mathbb{R}^m$. For this setting, $T = 2n$ is the current best estimate of the statistical threshold for phase retrieval [29]. The matrix X^\natural is set to $x^\natural(x^\natural)^\top$ where x^\natural is drawn uniformly from the surface




Figure 1: The convergence of the Polyak-variant of RMD on sparse recovery, low-rank matrix sensing, and phase retrieval for different values of m , the number of observations. The x -axis measures iterations in *thousands* of iterations and the y -axis is the distance between the true signal and the current iterate $\|X_t - X^\natural\|$ measured in the 1-norm for sparse recovery and in the nuclear norm for both low-rank matrix sensing and phase retrieval.

of the sphere. The linear operator \mathcal{A} is defined as $\mathcal{A}(X) = \text{diag}(GXG^\top)$ where $G \in \mathbb{R}^{m \times n}$ has independent Gaussian entries with variance $1/m$. We run **Polyak-RMD** on the function

$$\text{tr}(X) + 3 \|\mathcal{A}(X) - b\|_1 + 2 \text{dist}(X, \mathbb{S}_+^n).$$

8.2 Experiments and results

8.2.1 Convergence rate and sharpness vs. number of samples

Our first set of experiments investigates how the convergence rate of **Polyak-RMD** behaves as the number of samples is increased. We test sparse recovery with $(n, k) = (10000, 5)$ and $m \in \{T, 2T, 3T, 4T\}$, low-rank matrix sensing with $(n, k) = (100, 5)$ and $m \in \{T, 2T, 3T, 4T\}$, and phase retrieval with $n = 100$ and $m \in \{4T, 8T, 16T, 32T\}$.

For each of the three problems and for each choice of m , we compare the distance of $\|x_k - x^\natural\|$ against the iteration number k (see Figure 1). In all three settings, we see that the convergence rate of **Polyak-RMD** is linear (in terms of number of iterations) with a rate that improves as the number of samples increases. Table 2 indicates the maximum number of iterations of mirror descent performed in any round of **Polyak-RMD** in each of the three settings and for each choice of m . We believe this gives good insight into how the quantity μ varies with the number of samples m . Indeed, recall by Theorem 2 that the maximum number of iterations in any round of RMD (and hence **Polyak-RMD**) is *bounded* by $\lceil \frac{e^3 L^2 \ln(n)}{\mu^2} \rceil$. We indicate also the value of μ suggested by this formula—note that this may not be the true value of μ and is only an upper bound. Again, we see that the suggested value of μ improves (i.e., increases) with the number of observations in all three of our settings.

8.2.2 Convergence rate vs. dimension

Our second set of experiments investigates how the convergence rate of **Polyak-RMD** depends on the ambient dimension. We test sparse recovery with $n \in \{10^4, 10^5, 10^6\}$, $k = 5$, and $m \in \{T, 2T, 3T, 4T\}$. We compare the performance of **Polyak-RMD** with that of the subgradient method with Polyak

m	Max. iter.	μ	m	Max. iter.	μ	m	Max. iter.	μ
T	2.3×10^4	1.3	T	1.4×10^4	1.2	$4T$	4.2×10^3	1.3
$2T$	4.5×10^3	2.9	$2T$	2.4×10^3	2.8	$8T$	1.6×10^3	2.2
$3T$	1.1×10^3	5.9	$3T$	1.2×10^3	3.9	$16T$	7.6×10^2	3.1
$4T$	9.4×10^2	6.4	$4T$	9.0×10^2	4.6	$32T$	5.6×10^2	3.7

(a) Sparse recovery (b) Low-rank matrix sensing (c) Phase retrieval

Table 2: Experimental results for sparse recovery with $(n, k) = (100000, 5)$, low-rank matrix sensing with $(n, k) = (100, 5)$, and phase retrieval with $n = 100$. For each setting and each choice of m , we record the maximum number of iterations used to complete one round of Polyak-RMD. We additionally compute the value of μ suggested by the formula $\mu \approx \sqrt{e^3 L^2 \ln(n)/t}$ where t is the maximum number of iterations in any round of mirror descent.




Figure 2: The convergence of Polyak-RMD and Polyak-GD on sparse recovery for different values of n , the ambient dimension. The support size of x^\sharp is five. The x -axis measures iterations in thousands of iterations and the y -axis is the distance between the true signal and the current iterate $\|x_t - x^\sharp\|_1$ measured in the 1-norm.

step-sizes (Polyak-GD) [48]. The latter method also assumes that the value of f^* is known and is guaranteed to converge linearly on Euclidean sharp Lipschitz convex functions. See Figure 2 for the experimental results. As expected, the performance of Polyak-GD degrades with dimension as the sharpness of the function in question w.r.t. $\|\cdot\|_2$ should scale as $\approx n^{-1/2}$. In contrast, the performance of Polyak-RMD only varies minorly as the dimension n increases.

9 Discussion

This paper shows that for various statistical signal recovery tasks, once the sample size is greater than a constant multiple of the recovery threshold, then the convex optimization problem becomes well-conditioned in the sense of sharpness w.r.t. the ℓ_1 or Schatten-1 norm. In turn, this fact shows the optimization problem is robust to measurement error and optimization error. Furthermore, the newly developed algorithm RMD is able to achieve nearly-dimension-independent convergence rates.

We hope this paper induces interest in the interplay between statistics and convex optimization, especially in nonsmooth formulations and algorithms for these nonsmooth formulations. In particular, the following two directions might be of interest for future investigations:

- **Well conditioning beyond RIP:** This paper considers recovery tasks where strong RIP bounds have been established. Can we prove well-conditioning results for other statistical problems such as matrix completion [16] and phase retrieval with coded diffraction patterns [11] where similarly strong RIP bounds may not hold?
- **Adaptive penalty parameters:** Given a concrete statistical model, we can give upper bounds on r and ℓ so that the corresponding function $F_{r,\ell}$ is μ -sharp. On the other hand, determining these parameters may be difficult without a precise statistical model. Thus, are there algorithmic ways of estimating or adaptively choosing these two parameters?

Acknowledgments

Part of this research was completed in part while Alex L. Wang was supported by the Dutch Scientific Council (NWO) grant OCENW.GROOT.2019.015 (OPTIMAL). Lijun Ding was supported by an NSF TRIPODS grant to the Institute for Foundations of Data Science (NSF DMS-2023239) and by DOE ASCR Subcontract 8F-30039 from Argonne National Laboratory. Lijun Ding would like to thank Dmitriy Drusvyatskiy and Mateo Diaz for helpful discussions.

References

- [1] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton. Complementarity and nondegeneracy in semidefinite programming. *Mathematical programming*, 77(1):111–128, 1997.
- [2] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference: A Journal of the IMA*, 3(3):224–294, 2014.
- [3] H. Attouch, J. Bolte, P. Redont, and A. Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the kurdyka-łojasiewicz inequality. *Mathematics of operations research*, 35(2):438–457, 2010.
- [4] K. Ball, E. A. Carlen, and E. H. Lieb. Sharp uniform convexity and smoothness inequalities for trace norms. *Inequalities: Selecta of Elliott H. Lieb*, pages 171–190, 2002.
- [5] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive approximation*, 28:253–263, 2008.
- [6] H. H. Bauschke, J. M. Borwein, and W. Li. Strong conical hull intersection property, bounded linear regularity, jameson’s property (g), and error bounds in convex optimization. *Mathematical Programming*, 86:135–160, 1999.
- [7] H. H. Bauschke, P. L. Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- [8] S. Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

- [9] T. T. Cai and A. Zhang. Rop: Matrix recovery via rank-one projections. *The Annals of Statistics*, 43(1):102–138, 2015.
- [10] E. J. Candès and X. Li. Solving quadratic equations via phaselift when there are about as many equations as unknowns. *Foundations of Computational Mathematics*, 14(5):1017–1026, 2014.
- [11] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Applied and Computational Harmonic Analysis*, 39(2):277–299, 2015.
- [12] E. J. Candès and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2011.
- [13] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.
- [14] E. J. Candès, T. Strohmer, and V. Voroninski. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- [15] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215, 2005.
- [16] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.
- [17] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12:805–849, 2012.
- [18] V. Charisopoulos, Y. Chen, D. Davis, M. Díaz, L. Ding, and D. Drusvyatskiy. Low-rank matrix recovery with composite optimization: good conditioning and rapid convergence. *Foundations of Computational Mathematics*, 21(6):1505–1593, 2021.
- [19] V. Charisopoulos, D. Davis, M. Díaz, and D. Drusvyatskiy. Composite optimization for robust blind deconvolution. *arXiv preprint arXiv:1901.01624*, 2019.
- [20] Y. Chen. Incoherence-optimal matrix completion. *IEEE Transactions on Information Theory*, 61(5):2909–2923, 2015.
- [21] Y. Chen, Y. Chi, and A. J. Goldsmith. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Transactions on Information Theory*, 61(7):4034–4059, 2015.
- [22] D. Davis, D. Drusvyatskiy, K. J. MacPhee, and C. Paquette. Subgradient methods for sharp weakly convex functions. *Journal of Optimization Theory and Applications*, 179:962–982, 2018.
- [23] L. Ding and Y. Chen. Leave-one-out approach for matrix completion: Primal and dual analysis. *IEEE Transactions on Information Theory*, 66(11):7274–7301, 2020.
- [24] L. Ding and M. Udell. A strict complementarity approach to error bound and sensitivity of solution of conic programs. *Optimization Letters*, pages 1–24, 2022.

- [25] L. Ding, A. Yurtsever, V. Cevher, J. A. Tropp, and M. Udell. An optimal-storage approach to semidefinite programming using approximate complementarity. *SIAM Journal on Optimization*, 31(4):2695–2725, 2021.
- [26] D. Drusvyatskiy, A. D. Ioffe, and A. S. Lewis. Generic minimizing behavior in semialgebraic optimization. *SIAM Journal on Optimization*, 26(1):513–534, 2016.
- [27] D. Drusvyatskiy and A. S. Lewis. Error bounds, quadratic growth, and linear convergence of proximal methods. *Mathematics of Operations Research*, 43(3):919–948, 2018.
- [28] D. Drusvyatskiy and C. Paquette. Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming*, 178:503–558, 2019.
- [29] J. C. Duchi and F. Ruan. Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. *Information and Inference: A Journal of the IMA*, 8(3):471–529, 2019.
- [30] M. Dür, B. Jargalsaikhan, and G. Still. Genericity results in linear conic programming—a tour d’horizon. *Mathematics of operations research*, 42(1):77–94, 2017.
- [31] A. Edelman. Eigenvalues and condition numbers of random matrices. *SIAM journal on matrix analysis and applications*, 9(4):543–560, 1988.
- [32] I. I. Eremin. The relaxation method of solving systems of inequalities with convex functions on the left-hand side. In *Doklady Akademii Nauk*, volume 160, pages 994–996. Russian Academy of Sciences, 1965.
- [33] M. C. Ferris. Weak sharp minima and exact penalty functions. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 1988.
- [34] D. Garber and A. Kaplan. On the efficient implementation of the matrix exponentiated gradient algorithm for low-rank matrix optimization. *Mathematics of Operations Research*, 2022.
- [35] J.-L. Goffin. On convergence rates of subgradient optimization methods. *Mathematical programming*, 13:329–347, 1977.
- [36] D. Goldfarb and K. Scheinberg. Interior point trajectories in semidefinite programming. *SIAM Journal on Optimization*, 8(4):871–886, 1998.
- [37] P. Hand. Conditions for existence of dual certificates in rank-one semidefinite problems. *arXiv*, 1303.1598, 2013.
- [38] A. J. Hoffman. On approximate solutions of systems of linear inequalities. In *Selected Papers Of Alan J Hoffman: With Commentary*, pages 174–176. World Scientific, 2003.
- [39] P. R. Johnstone and P. Moulin. Faster subgradient methods for functions with hölderian growth. *Mathematical Programming*, 180(1-2):417–450, 2020.
- [40] A. S. Lewis. Nonsmooth optimization: conditioning, convergence and semi-algebraic models. In *Proceedings of the International Congress of Mathematicians, Seoul*, volume 4, pages 872–895, 2014.
- [41] X. Li, Z. Zhu, A. Man-Cho So, and R. Vidal. Nonconvex robust low-rank matrix recovery. *SIAM Journal on Optimization*, 30(1):660–686, 2020.

- [42] M. V. Nayakkankuppam and M. L. Overton. Conditioning of semidefinite programs. *Mathematical programming*, 85(3), 1999.
- [43] I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *Mathematical Programming*, 175:69–107, 2019.
- [44] A. S. Nemirovski and Y. E. Nesterov. Optimal methods of smooth convex minimization. *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki*, 25(3):356–369, 1985.
- [45] Y. Nesterov and A. Nemirovski. On first-order algorithms for ℓ_1 /nuclear norm minimization. *Acta Numerica*, 22:509–575, 2013.
- [46] G. Pataki and L. Tunçel. On the generic properties of convex optimization problems in conic form. *Mathematical Programming*, 89:449–457, 2001.
- [47] Y. Plan and R. Vershynin. Dimension reduction by random hyperplane tessellations. *Discrete & Computational Geometry*, 51(2):438–461, 2014.
- [48] B. T. Polyak. Minimization of unsmooth functionals. *USSR Computational Mathematics and Mathematical Physics*, 9(3):14–29, 1969.
- [49] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(12), 2011.
- [50] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- [51] J. Renegar and B. Grimmer. A simple nearly optimal restart scheme for speeding up first-order methods. *Foundations of Computational Mathematics*, 22(1):211–256, 2022.
- [52] R. T. Rockafellar. *Convex analysis*, volume 11. Princeton university press, 1997.
- [53] V. Roulet, N. Boumal, and A. d’Aspremont. Renegar’s condition number, sharpness and compressed sensing performance. *arXiv preprint arXiv:1506.03295*, 2015.
- [54] V. Roulet and A. d’Aspremont. Sharpness, restart and acceleration. *Advances in Neural Information Processing Systems*, 30, 2017.
- [55] M. Rudelson and R. Vershynin. Sparse reconstruction by convex relaxation: Fourier and gaussian measurements. In *2006 40th Annual Conference on Information Sciences and Systems*, pages 207–212. IEEE, 2006.
- [56] J. F. Sturm. Error bounds for linear matrix inequalities. *SIAM Journal on Optimization*, 10(4):1228–1248, 2000.
- [57] R. Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, page 210–268. Cambridge University Press, 2012.
- [58] A. L. Wang and F. Kilinc-Karzan. Accelerated first-order methods for a class of semidefinite programs. *arXiv preprint arXiv:2206.00224*, 2022.
- [59] T. Yang and Q. Lin. Rsg: Beating subgradient method without smoothness and strong convexity. *The Journal of Machine Learning Research*, 19(1):236–268, 2018.

- [60] C. Zalinescu. *Convex analysis in general vector spaces*. World scientific, 2002.
- [61] S. Zhang. Global error bounds for convex conic problems. *SIAM Journal on Optimization*, 10(3):836–851, 2000.
- [62] Z. Zhou and A. M.-C. So. A unified approach to error bounds for structured convex optimization problems. *Mathematical Programming*, 165:689–728, 2017.

A Proof of Lemma 2

Proof. The forward direction follows directly from Lemma 1. We focus on the reverse direction. For notational convenience, let $\mathcal{S}_1 := \partial f(x^\natural)$, $\mathcal{S}_2 := -\text{range}(\mathcal{A}^*)$ and $\mathcal{S}_3 := -(K^* \cap (x^\natural)^\perp)$. The inclusion (7) implies the right of (7) is full-dimensional. Thus, its interior coincides with its relative interior. By additivity of the relative interior [7, Corollary 6.15], we have

$$0 \in \text{rint}(\mathcal{S}_1) + \text{rint}(\mathcal{S}_2) + \text{rint}(\mathcal{S}_3).$$

We deduce there exists $g^\natural \in \text{rint}(\mathcal{S}_1)$, $h^\natural \in \text{rint}(\mathcal{S}_2)$, and $z^\natural \in \text{rint}(\mathcal{S}_3)$ such that

$$0 = g^\natural + h^\natural + z^\natural.$$

There exists an $\epsilon > 0$, such that

$$\begin{aligned} B_{V^*}(g^\natural, \epsilon) \cap \text{aff}(\mathcal{S}_1) &\subseteq \mathcal{S}_1, \\ B_{V^*}(h^\natural, \epsilon) \cap \text{aff}(\mathcal{S}_2) &\subseteq \mathcal{S}_2, \quad \text{and} \\ B_{V^*}(z^\natural, \epsilon) \cap \text{aff}(\mathcal{S}_3) &\subseteq \mathcal{S}_3. \end{aligned}$$

Let $r > 0$ such that $-\mathcal{A}^*(B_{W^*}(0, r)) \supseteq B_{V^*}(h^\natural, \epsilon) \cap \text{aff}(\mathcal{S}_2)$ and let $\ell = \|z^\natural\| + \epsilon$. Then,

$$\begin{aligned} &\dim \left(\partial f(x^\natural) - \mathcal{A}^*(B_{W^*}(0, r)) - (K^* \cap (x^\natural)^\perp \cap B_{V^*}(0, \ell)) \right) \\ &\stackrel{(a)}{\geq} \dim \left(B_{V^*}(g^\natural, \epsilon) \cap \mathcal{S}_1 + B_{V^*}(h^\natural, \epsilon) \cap \text{aff}(\mathcal{S}_2) - B_{V^*}(z^\natural, \epsilon) \cap \text{aff}(\mathcal{S}_3) \right) \\ &\stackrel{(b)}{=} \dim(\text{aff}(\mathcal{S}_1 + \mathcal{S}_2 + \mathcal{S}_3)) \\ &\stackrel{(c)}{=} \dim(V^*). \end{aligned}$$

Here, step (a) follows as dimension is set-monotone, step (b) follows as the dimension of a convex set is defined as the dimension of its affine hull, and the last step (d) follows by the assumption that $0 \in \text{int}(\mathcal{S}_1 + \mathcal{S}_2 + \mathcal{S}_3)$. We deduce that $\partial f(x^\natural) - \mathcal{A}^*(B_{W^*}(0, r)) - (K^* \cap (x^\natural)^\perp \cap B_{V^*}(0, \ell))$ is a full-dimensional convex set containing 0 in its relative interior. Thus, it contains $B_{V^*}(0, \mu)$ for some $\mu > 0$. \blacksquare

B Proof of Proposition 1 for sparse vector recovery

Proof. Let $\delta \in \mathbb{R}^n$ be arbitrary. Our goal is to show that

$$\|x^\natural + \delta\|_1 + \sqrt{k'} \|A\delta\|_1 - \|x^\natural\| \geq \frac{\epsilon}{2 + \epsilon} \|\delta\|.$$

Without loss of generality, we may reindex \mathbb{R}^n so that $\text{supp}(x^\natural) \subseteq [k]$ and $|\delta_{k+1}| \geq |\delta_{k+2}| \geq \dots \geq |\delta_n|$. We decompose δ as a sum of vectors $\delta_k + \sigma_1 + \dots + \sigma_t$, each in \mathbb{R}^n . Specifically, δ_k extracts the first k coordinates of δ , σ_1 extracts the next k' coordinates of δ , and each subsequent σ_i extracts the next k' coordinates of δ . Finally, σ_t may extract fewer than k' coordinates of δ . Let $\delta_{k^\perp} := \delta - \delta_k = \sigma_1 + \dots + \sigma_t$.

We can bound

$$\begin{aligned}
\|\delta_{k^\perp}\|_1 &\geq \sum_{i=1}^{t-1} \|\sigma_i\|_1 && (\sigma_i \text{ are disjoint}) \\
&= k' \sum_{i=1}^{t-1} \frac{\|\sigma_i\|_1}{k'} \\
&\geq k' \sum_{i=2}^t \|\sigma_i\|_\infty && (|\delta_{k+1}| \geq \dots \geq |\delta_n|) \\
&\geq \sqrt{k'} \sum_{i=2}^t \|\sigma_i\|_2 \\
&\geq \sqrt{k'} \cdot (\text{RIP}_{k'}^+)^{-1} \sum_{i=2}^t \|A\sigma_i\|_1 \\
&\geq \sqrt{k'} \cdot (\text{RIP}_{k'}^+)^{-1} (\|A(\delta_k + \sigma_1)\|_1 - \|A\delta\|_1) && (\text{triangle inequality}) \\
&\geq \sqrt{k'} \frac{\text{RIP}_{k+k'}^-}{\text{RIP}_{k'}^+} \|\delta_k + \sigma_1\|_2 - \sqrt{k'} \|A\delta\|_1 \\
&\geq \sqrt{k'} \frac{\text{RIP}_{k+k'}^-}{\text{RIP}_{k'}^+} \|\delta_k\|_2 - \sqrt{k'} \|A\delta\|_1 && (\delta_k, \sigma_1 \text{ are disjoint}) \\
&\geq \sqrt{\frac{k'}{k}} \frac{\text{RIP}_{k+k'}^-}{\text{RIP}_{k'}^+} \|\delta_k\|_1 - \sqrt{k'} \|A\delta\|_1 \\
&\geq (1 + \epsilon) \|\delta_k\|_1 - \sqrt{k'} \|A\delta\|_1.
\end{aligned}$$

We are now ready to prove sharpness:

$$\begin{aligned}
&\left\| x^\natural + \delta \right\|_1 + \sqrt{k'} \|A\delta\|_1 - \left\| x^\natural \right\|_1 \\
&\geq \|\delta_{k^\perp}\|_1 - \|\delta_k\|_1 + \sqrt{k'} \|A\delta\|_1 && (x^\natural + \delta_k, \delta_{k^\perp} \text{ are disjoint}) \\
&= \left(\frac{2}{2+\epsilon} + \frac{\epsilon}{2+\epsilon} \right) \|\delta_{k^\perp}\|_1 - \|\delta_k\|_1 + \sqrt{k'} \|A\delta\|_1 \\
&\geq \left(\frac{2}{2+\epsilon} (1+\epsilon) - 1 \right) \|\delta_k\|_1 + \frac{\epsilon}{2+\epsilon} \|\delta_{k^\perp}\|_1 && (\text{bound on } \|\delta_{k^\perp}\|_1) \\
&\quad + \left(1 - \frac{2}{2+\epsilon} \right) \sqrt{k'} \|A\delta\|_1 \\
&\geq \frac{\epsilon}{2+\epsilon} \|\delta\|_1.
\end{aligned}$$

This proves the claim as $\delta \in \mathbb{R}^n$ was arbitrary. ■