



HAL
open science

Beyond Moore-Penrose Part II: The Sparse Pseudoinverse

Ivan Dokmanić, Rémi Gribonval

► **To cite this version:**

Ivan Dokmanić, Rémi Gribonval. Beyond Moore-Penrose Part II: The Sparse Pseudoinverse. 2017. hal-01547283v1

HAL Id: hal-01547283

<https://inria.hal.science/hal-01547283v1>

Preprint submitted on 30 Jun 2017 (v1), last revised 13 Jul 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Beyond Moore-Penrose

Part II: The Sparse Pseudoinverse

Ivan Dokmanić and Rémi Gribonval

Abstract

This is the second part of a two-paper series on generalized inverses that minimize matrix norms. In Part II we focus on generalized inverses that are minimizers of entrywise ℓ^p norms whose main representative is the *sparse pseudoinverse* for $p = 1$. We are motivated by the idea to replace the Moore-Penrose pseudoinverse by a sparser generalized inverse which is in some sense well-behaved. Sparsity implies that it is faster to apply the resulting matrix; well-behavedness would imply that we do not lose much in stability with respect to the least-squares performance of the MPP.

We first address questions of uniqueness and non-zero count of (putative) sparse pseudoinverses. We show that a sparse pseudoinverse is generically unique, and that it indeed reaches optimal sparsity for almost all matrices. We then turn to proving our main stability result: finite-size concentration bounds for the Frobenius norm of ℓ^p -minimal inverses for $1 \leq p \leq 2$. Our proof is based on tools from convex analysis and random matrix theory, in particular the recently developed convex Gaussian min-max theorem. Along the way we prove several results about sparse representations and convex programming that were known folklore, but of which we could find no proof.

1 Introduction

A generalized inverse (or a pseudoinverse) of a rectangular matrix is a matrix that has some properties of the usual inverse of a regular square matrix. We call $\mathbf{X} \in \mathbb{C}^{n \times m}$ a generalized inverse of $\mathbf{A} \in \mathbb{C}^{m \times n}$ if it holds that $\mathbf{A}\mathbf{X}\mathbf{A} = \mathbf{A}$; we denote the set of all generalized inverses of \mathbf{A} by $\mathcal{G}(\mathbf{A})$.

The most common generalized inverse is the famous Moore-Penrose pseudoinverse (MPP). Some of the many uses of the MPP are to compute the least-squares fit to an overdetermined system of linear equations, or to find the shortest solution of an underdetermined system; many other related results are presented in Part I.

However, the MPP is only one out of infinitely many possible generalized inverses. By relaxing the requirements needed to get the MPP we free up degrees of freedom that can be optimized to promote other interesting properties. Our study focuses on a particular strategy to obtain alternative pseudoinverses: among all pseudoinverses, the MPP has the smallest Frobenius norm¹, so it seems natural to ask what happens if we replace the Frobenius norm by other matrix norms.

A generalized inverse of $\mathbf{A} \in \mathbb{C}^{m \times n}$, $m < n$ that minimizes some arbitrary (quasi)norm ν is defined as

$$\text{ginv}_\nu(\mathbf{A}) \stackrel{\text{def}}{=} \arg \min_{\mathbf{X}} \|\mathbf{X}\|_\nu \quad \text{subject to } \mathbf{X} \in \mathcal{G}(\mathbf{A}).$$

¹It also minimizes a number of other norms (see Part I).

Strictly speaking, $\text{ginv}_\nu(\mathbf{A})$ is a set, since the corresponding optimization program may have more than one solution. This, however, will not be the case for most studied norms and matrices (see, for example, Section 2) so we write “=” instead of “ \in ” and warn the reader when extra care is advised. We also note that if $\|\cdot\|_\nu$ is a bona fide norm, then $\text{ginv}_\nu(\mathbf{A})$ involves solving a convex program. At least in principle, this means that it can actually be computed in the sense that any first-order scheme leads to a global optimum.

In this Part II, we specifically focus on entrywise ℓ^p norms² which are straightforward extensions of vector ℓ^p norms. Concretely, for $\mathbf{M} \in \mathbb{C}^{m \times n}$ and $0 < p < \infty$ we have $\|\mathbf{M}\|_p \stackrel{\text{def}}{=} \|\text{vec}(\mathbf{M})\|_p = \left(\sum_{ij} |m_{ij}|^p\right)^{1/p}$, with $\text{vec}(\cdot)$ denoting the concatenation of the columns of the argument matrix. A particular entrywise norm is the Frobenius norm associated to $p = 2$ and minimizing it gives the MPP. But our main motivation to look at entrywise norms is rather the case $p = 1$, which, as we show, leads to sparse pseudoinverses.

The motivation is that applying a sparse pseudoinverse requires less computation than applying a full one [DKV13, LLM13, CHKK, KKL12]. We could take advantage of this fact if we knew how to compute a sparse pseudoinverse that is in some sense stable to noise. Ignoring the last requirement, finding the sparsest pseudoinverse may be formulated as

$$\text{ginv}_0(\mathbf{A}) \stackrel{\text{def}}{=} \arg \min_{\mathbf{X}} \|\text{vec}(\mathbf{X})\|_0 \text{ subject to } \mathbf{X} \in \mathcal{G}(\mathbf{A}), \quad (1)$$

where $\|\cdot\|_0$ counts the total number of non-zero entries in a vector and $\text{vec}(\cdot)$ transforms a matrix into a vector by stacking its columns. The non-zero count gives the naive complexity of applying \mathbf{X} or its adjoint to a vector (for an illustration, see Figure 1).

Any optimally sparse generalized inverse of \mathbf{A} in the sense of the ℓ^0 norm (1) is by definition in the set $\text{ginv}_0(\mathbf{A})$. When $\mathbf{A} \in \mathbb{C}^{m \times n}$, $m < n$, has full column rank the condition $\mathbf{X} \in \mathcal{G}(\mathbf{A})$ is equivalent to $\mathbf{A}\mathbf{X} = \mathbf{I}_m$, and computing an element from this set can be expressed columnwise, as a collection of ℓ^0 minimization problems

$$[\text{ginv}_0(\mathbf{A})]_i = \arg \min_{\mathbf{A}\mathbf{x}=\mathbf{e}_i} \|\mathbf{x}\|_{\ell^0}.$$

Even though optimization problems of this kind are NP-hard [Nat95, DMA97], we will see that for most matrices \mathbf{A} finding a solution is trivial and not very useful: just invert any full-rank $m \times m$ submatrix and set the rest to zero. What we mean by “not useful” is that the resulting matrix is in general poorly conditioned. On the other hand, the vast literature on conditions of equivalence between solutions of ℓ^0 and ℓ^1 minimization (see, e.g. [Don06]) suggests that ℓ^1 minimization is a good computationally tractable proxy. In fact, we will show that the sparsest pseudoinverse can be computed by convex optimization; that is, a convex program generically provides a minimizer of the ℓ^0 norm. Moreover, we will show that this minimizer is generically unique. This motivates the definition of the specific notation

$$\text{spinv}(\mathbf{A}) \stackrel{\text{def}}{=} \text{ginv}_1(\mathbf{A}) = \arg \min_{\mathbf{X}} \|\text{vec}(\mathbf{X})\|_1 \text{ subject to } \mathbf{X} \in \mathcal{G}(\mathbf{A}). \quad (2)$$

Not only is (2) computationally tractable but we will show that it leads to well-behaved matrices that are indeed sparse. Intuitive reasoning is as follows: an $m \times n$ matrix \mathbf{A} is generically full rank, hence the constraint $\mathbf{X} \in \mathcal{G}(\mathbf{A})$ is generically equivalent to $\mathbf{A}\mathbf{X} = \mathbf{I}_m$. As this is a shorthand notation for m^2 linear equations, it constrains m^2 degrees of freedom. The matrix \mathbf{X} has nm entries, leaving us with $nm - m^2$ degrees of freedom, which will (hopefully) be set to zero by ℓ^1 minimization.

²For brevity, we loosely call “norm” any quasi-norm such as ℓ^p , $p < 1$, as well as the “pseudo-norm” ℓ^0 .




Figure 1: Illustration of the sparse pseudoinverse: \mathbf{A} is a realization of a 10×30 iid Gaussian matrix; the MPP $\mathbf{X}_1 = \mathbf{A}^\dagger$ is a full matrix; the sparse pseudoinverse $\mathbf{X}_2 = \text{spinv}(\mathbf{A})$ has only 10^2 non-zeros among 10×30 entries.

1.1 Prior Art

Several earlier works address sparsity of generalized inverses and dual frames [CHKK, KKL12, LLM13]. These works concentrate on existence results and explicit constructions of sparse frames and sparse dual frames with prescribed spectra. Krahmer, Kutyniok, and Lemvig [KKL12] establish sharp bounds on the sparsity of dual frames, showing that generically, for $\mathbf{A} \in \mathbb{C}^{m \times n}$, the sparsest dual has $mn - m^2$ zeros. Li, Liu, and Mi [LLM13] provide bounds on the sparsity of duals of Gabor frames which are better than generic bounds. They also introduce the idea of using ℓ^p (for $p = 1$) minimization to compute these dual frames, and show that under certain conditions, the ℓ^1 minimization yields the sparsest possible dual Gabor frame. Further examples of non-canonical dual Gabor frames are given by Perraudin *et al.*, who use convex optimization to derive dual Gabor frames with more favorable properties than the canonical one [PHSB14], particularly in terms of time-frequency localization.

Our stability results fit squarely within random matrix theory. However, a perusal of the corresponding rich literature indicates that results on finite-size concentration bounds for norms of pseudoinverses are scarce. One notable exception is [HMT11] which gives the expected Frobenius norm of the MPP and a corresponding upper bound on a probability of large deviation; we note that our techniques allow us to get a finite-size concentration result for the MPP for a range of deviations complementary to the one considered in [HMT11]. It is also worth noting that there are a number of results for square matrices [Rud08, Ver14].

Finally, the sparse pseudoinverse was previously studied in [DKV13], where it was shown empirically that the minimizer is indeed a sparse matrix, and that it can be used to speed up the resolution of certain inverse problems.

1.2 Our Contributions and Paper Outline

We address fundamental questions of uniqueness and stability of the *sparse pseudoinverse*. First, in Section 2, we show that minimizing the entrywise ℓ^1 norm of the generalized inverse generically results in a maximally sparse inverse. By exploiting connections between basis pursuit and basis pursuit denoising/lasso we also show that the optimum is generically unique. We then prove

in Section 3 that unlike other, simpler strategies that yield a sparsest generalized inverse, ℓ^1 minimization produces a *good*, well-conditioned matrix.

We measure well-behavedness by the Frobenius norm, which arises naturally in the context of linear inverse problems with noisy measurements: for white noise the output MSE is directly proportional to the Frobenius norm. Our main result is a characterization of the typical Frobenius norm of the ℓ^p -minimal generalized inverse for random Gaussian matrices for $1 \leq p \leq 2$. As a corollary we show the the Frobenius norm of the sparse pseudoinverse $\text{ginv}_1(\mathbf{A})$ is indeed controlled.

The proof uses the *convex Gaussian min-max theorem* [TAH16, OTH13, TO15] which has thus far found major use in quantifying the performance of regularized M-estimators such as the lasso [TAH16]. Unlike these previous applications, we give finite-size concentration bounds rather than “in probability” asymptotic results.

Along the way we collect a number of useful byproducts related to matrix analysis, sparse representations, and convex programming. For example, we prove that basis pursuit generically has a unique minimizer—a folklore fact mentioned repeatedly in the literature, but of which we could find no proof. The same goes for the fact that ℓ_p minimization for $0 \leq p \leq 1$ has a sparse solution which is to the best of our knowledge yet another piece of thus far unproven folklore.

2 A Sparse Pseudoinverse is (Generically) Unique and Optimally Sparse

As the section title promises, we first settle the questions of uniqueness and non-zero count of a “sparse” pseudoinverse $\text{spinv}(\mathbf{A})$. We rely on two groups of facts: first, existing results on sparse dual frames which characterize maximal levels of sparsity attainable by generalized inverses [KKL12, CHKK, LLM13], and second, new results on constrained ℓ^p minimization which we put forward in Section 2.1.

Our main result in this section can be informally stated as

Theorem 2.1 (Informal). *For a generic $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m < n$, $\text{spinv}(\mathbf{A})$ is unique, $\text{spinv}(\mathbf{A}) \in \text{ginv}_0(\mathbf{A})$, and it has exactly m non-zeros per column.*

We begin by proving two preliminary results about vectors.

2.1 Vector Facts

First, we show that the set of ℓ^p minimizers for $0 \leq p \leq 1$ always contains an m -sparse point. While this is known folklore, we could only find a proof for $p = 1$ [Ela10, Section 1.4] so we give one for $0 \leq p \leq 1$ below, courtesy of Simon Foucart:

Lemma 2.1. *Consider $0 \leq p \leq 1$ and let $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m < n$ and $\mathbf{y} \in \mathcal{R}(\mathbf{A})$. Then the set of minimizers of $\|\mathbf{z}\|_p$ subject to the constraint that $\mathbf{Az} = \mathbf{y}$ always contains a point with at most m non-zero entries.*

Proof. Let $\mathcal{Z}(\mathbf{y})$ be the (non-empty, closed) set of all solutions to $\mathbf{y} = \mathbf{Az}$ that reach the minimum ℓ^p norm solution. Let $\mathbf{z} \in \mathcal{Z}(\mathbf{y})$ be an element of this set reaching the smallest ℓ^0 norm: $\|\mathbf{z}\|_0 \leq \|\mathbf{z}'\|_0$ for all $\mathbf{z}' \in \mathcal{Z}(\mathbf{y})$. For the sake of contradiction, assume that $\|\mathbf{z}\|_0 > m$ and denote by S the support of \mathbf{z} . By assumption the size of S exceeds m , the number of rows of \mathbf{A} , hence the columns of \mathbf{A} indexed by S are linearly dependent and there exists a nonzero vector \mathbf{h} of the null space of \mathbf{A} that is supported in S . Considering $\mathbf{z}' = \mathbf{z} + \epsilon \mathbf{h}$ we have $\mathbf{Az}' = \mathbf{y}$.

There is a critical $\epsilon_o > 0$ and $s \in \{\pm 1\}$ such that: $\text{sign}(\mathbf{z} + \epsilon \mathbf{h}) = \text{sign}(\mathbf{z})$ when $|\epsilon| < \epsilon_o$; and $\|\mathbf{z} + s\epsilon_o \mathbf{h}\|_0 < \|\mathbf{z}\|_0$ (corresponding to the cancellation of at least one component of \mathbf{z} by adding $s\epsilon_o \mathbf{h}$). In the remainder of the proof we take $s = 1$ with no loss of generality.

For $p = 0$, the contradiction follows immediately. For $0 < p \leq 1$, we let $q \stackrel{\text{def}}{=} 1/p$ and $q^* = 1/(1-p)$ so that $1/q + 1/q^* = 1$. Then by Hölder's inequality, when $|\epsilon| < \epsilon_o$

$$\begin{aligned} \|\mathbf{z}'\|_p^p &= \sum_{j \in S} \frac{|z'_j|^p}{|z_j|^{p(1-p)}} \cdot |z_j|^{p(1-p)} \leq \left(\sum_{j \in S} \left(\frac{|z'_j|^p}{|z_j|^{p(1-p)}} \right)^q \right)^{1/q} \cdot \left(\sum_{j \in S} (|z_j|^{p(1-p)})^{q^*} \right)^{1/q^*} \\ &= \left(\sum_{j \in S} \frac{|z'_j|^p}{|z_j|^{p(1-p)}} \right)^p \cdot \left(\sum_{j \in S} |z_j|^{p(1-p)} \right)^{1-p} = \left(\sum_{j \in S} \text{sign}(z_j + \epsilon h_j) \frac{z_j + \epsilon h_j}{|z_j|^{1-p}} \right)^p \cdot \|\mathbf{z}\|_p^{p(1-p)} \\ &= \left(\sum_{j \in S} \text{sign}(z_j) \frac{z_j + \epsilon h_j}{|z_j|^{1-p}} \right)^p \cdot \|\mathbf{z}\|_p^{p(1-p)} = \left(\|\mathbf{z}\|_p^p + \underbrace{\epsilon \sum_{j \in S} \text{sign}(z_j) \frac{h_j}{|z_j|^{1-p}}}_{\stackrel{\text{def}}{=} \sigma} \right)^p \cdot \|\mathbf{z}\|_p^{p(1-p)} \end{aligned}$$

Hence

- if $\sigma \neq 0$, choosing $|\epsilon| < \epsilon_o$ with $\text{sign}(\epsilon) = -\text{sign}(\sigma)$ we obtain $\|\mathbf{z}'\|_p^p < \|\mathbf{z}\|_p^p$; a contradiction.
- if $\sigma = 0$, choosing $\epsilon = s\epsilon_o$, by continuity we get $\|\mathbf{z}'\|_p^p \leq \|\mathbf{z}\|_p^p$ (hence $\mathbf{z}' \in \mathcal{Z}(\mathbf{y})$), yet $\|\mathbf{z}'\|_0 < \|\mathbf{z}\|_0$; another contradiction.

□

For the second fact, related to uniqueness, we need the notion of general position. To this end we quote two definitions from [Tib13].

Definition 2.1. *The $m \times n$ matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ has columns in general position if for any $k < m$ the affine span of any $k + 1$ vectors $\pm \mathbf{a}_{i_1}, \dots, \pm \mathbf{a}_{i_{k+1}}$ does not contain any vector $\pm \mathbf{a}_i$, $i \notin \{i_1, \dots, i_{k+1}\}$.*

Definition 2.2. *The $m \times n$ matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is in general position with respect to the vector $\mathbf{y} \in \mathbb{R}^m$ if for any $k < m$, \mathbf{y} does not belong to the linear span of any k columns of \mathbf{A} .*

We now show that for a generic $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{y} \in \mathbb{R}^m$, the problem

$$\min_{\mathbf{x}: \mathbf{A}\mathbf{x}=\mathbf{y}} \|\mathbf{x}\|_1 \tag{3}$$

has a unique minimizer. Similarly to Lemma 2.1, this fact seems to be a piece of folklore. But although it has been mentioned in the literature (e.g. [Dos12]) we could find no proof, so we produce one based on the properties of the lasso.

Lemma 2.2. *Assume that $\mathbf{A} \in \mathbb{R}^{m \times n}$ has columns in general position, and that \mathbf{A} is in general position with respect to $\mathbf{y} \in \mathbb{R}^m$. Then (3) has a unique solution.*

Proof of Lemma 2.2. The proof is a combination of a sufficient condition for uniqueness of the squared lasso [Fuc04, Tib13]:

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1 \quad (4)$$

with a technique to map solutions of (3) to those of (4) for small λ [Fuc04]. We start by writing out the KKT optimality conditions for the problem (4):

$$\mathbf{A}^\top(\mathbf{y} - \mathbf{A}\mathbf{x}^*) = \lambda \mathbf{h}, \text{ where } h_i \in \begin{cases} \{\text{sign}(x_i^*)\}, & x_i^* \neq 0, \\ [-1, 1], & x_i^* = 0 \end{cases}. \quad (5)$$

Next, the so-called equicorrelation set $I = I(\mathbf{y}, \mathbf{A}, \lambda)$ and equicorrelation sign $\mathbf{s} = \mathbf{s}(\mathbf{y}, \mathbf{A}, \lambda)$ are defined as

$$I \stackrel{\text{def}}{=} \{i : |\mathbf{a}_i^\top(\mathbf{y} - \mathbf{A}\mathbf{x}^*)| = \lambda\}, \quad \mathbf{s} \stackrel{\text{def}}{=} \text{sign}(\mathbf{A}_I^\top(\mathbf{y} - \mathbf{A}\mathbf{x}^*)). \quad (6)$$

The set I and the vector \mathbf{s} are unique in the sense that all solutions \mathbf{x}^* of (4) have the same equicorrelation set and equicorrelation sign. This follows from the strict convexity of the quadratic loss which, combined with the convexity of the ℓ^1 norm implies (see, e.g., [Tib13, Lemma 1]) that all solutions lead to the same $\mathbf{A}\mathbf{x}^*$. From these definitions and (5) we see that the solutions \mathbf{x}^* of (4) are characterized by

$$\mathbf{x}_{J^c}^* = \mathbf{0}, \quad (7)$$

$$\mathbf{A}_I^\top(\mathbf{y} - \mathbf{A}_I\mathbf{x}_I^*) = \lambda \mathbf{s}. \quad (8)$$

Since \mathbf{A} has columns in general position, as shown in [Tib13, Lemma 3], the matrix \mathbf{A}_I has a trivial nullspace for any \mathbf{y} and $\lambda > 0$, hence

$$|I| \leq m. \quad (9)$$

It follows that any solution \mathbf{x}^* of (4) can be written as

$$\mathbf{x}_I^* = \mathbf{A}_I^\dagger[\mathbf{y} - \lambda(\mathbf{A}_I^\top)^\dagger \mathbf{s}], \quad (10)$$

and the solution is indeed unique, with $\mathbf{s} = \text{sign}(\mathbf{x}_I^*)$.

Note that the solution (10) of (4) in general does not satisfy $\mathbf{A}\mathbf{x}^* = \mathbf{y}$. We now use a construction by Fuchs [Fuc04, Theorem 2] to relate solutions of (4) to those of (3).

Denote \mathcal{X} the convex set of solutions of (3). It is classical exercise in convex analysis [Ela10, Section 1.4] to show that this set contains at least one vector whose support J is associated to linearly independent columns of \mathbf{A} , i.e., \mathbf{A}_J has a trivial null space (see also Lemma 2.1 applied to $p = 1$). Hence, the set

$$\mathcal{J} \stackrel{\text{def}}{=} \{J = \text{support}(\mathbf{x}^\#) : \mathbf{x}^\# \in \mathcal{X}, \mathcal{N}(\mathbf{A}_J) = \{0\}\}$$

is non empty. Moreover, if (3) has a non-unique set of minimizers (i.e., if \mathcal{X} is *not* a singleton), then \mathcal{J} contains at least *two* distinct sets $J_1 \neq J_2$. To see this, first observe that \mathcal{X} is convex and compact so it is the convex hull of its extreme points. Let \mathbf{x}_e be an element of \mathcal{X} with support J , and assume \mathbf{A}_J has a non-trivial nullspace. Similarly as in Lemma 2.1, there exists a vector $\mathbf{h} \in \mathcal{N}(\mathbf{A})$ entirely supported in J . Then, for any ϵ , $\mathbf{x}_e + \epsilon \mathbf{h}$ is a feasible point for the optimization (3) and there exists a critical $\epsilon_0 > 0$ such that for all $|\epsilon| < \epsilon_0$, $\text{sign}(\mathbf{x}_e + \epsilon \mathbf{h}) = \text{sign}(\mathbf{x}_e) \stackrel{\text{def}}{=} \mathbf{s}$. For $|\epsilon| < \epsilon_0$, reasoning as in Lemma 2.1 we can compute $\|\mathbf{x}_e + \epsilon \mathbf{h}\|_1 = \|\mathbf{x}_e\|_1 + \epsilon \sum_{i=1}^n s_i h_i$. If

$\sigma \stackrel{\text{def}}{=} \sum_{i=1}^n s_i h_i \neq 0$ then by setting $\epsilon = -\text{sign}(\sigma)\epsilon_0/2$ we get $\|\mathbf{x}_e + \epsilon\mathbf{h}\|_1 = \|\mathbf{x}_e\|_1 - \frac{1}{2}\epsilon_0|\sigma| < \|\mathbf{x}_e\|_1$ which is a contradiction since $\mathbf{x}_e + \epsilon\mathbf{h}$ is feasible. It therefore must hold that $\sigma = 0$, meaning that $\|\mathbf{x}_e + \epsilon\mathbf{h}\|_1 = \|\mathbf{x}_e - \epsilon\mathbf{h}\|_1 = \|\mathbf{x}_e\|_1$ when $|\epsilon| < \epsilon_0$, or in other words, $(\mathbf{x}_e \pm \epsilon\mathbf{h}) \in \mathcal{X}$. As

$$\mathbf{x}_e = \frac{1}{2}(\mathbf{x}_e + \epsilon\mathbf{h}) + \frac{1}{2}(\mathbf{x}_e - \epsilon\mathbf{h}),$$

the element \mathbf{x}_e can be written as a convex combination (with non-zero coefficients) of two distinct points in \mathcal{X} which shows that it is not an extreme point of \mathcal{K} . By contraposition, this establishes that any extreme point \mathbf{x} of \mathcal{X} has a support $J(\mathbf{x})$ such that $\mathbf{A}_{J(\mathbf{x})}$ has full column rank, that is, $J(\mathbf{x}) \in \mathcal{J}$.

If \mathcal{J} is a singleton with one element J then \mathcal{X} has a single extreme point with non-zero entries $\mathbf{A}_J^\dagger \mathbf{y}$. Consequently \mathcal{X} is a singleton as the convex hull of that single extreme point. By contraposition, if \mathcal{X} is not a singleton, it has multiple extreme points and \mathcal{J} is not a singleton.

Consider any $J \in \mathcal{J}$ and \mathbf{x}^\sharp the corresponding solution of (3), characterized by $J = \text{support}(\mathbf{x}^\sharp)$, $\mathbf{x}_J^\sharp = \mathbf{A}_J^\dagger \mathbf{y}$ and $\mathbf{x}_{J^c}^\sharp = \mathbf{0}$. For any $\lambda > 0$ we construct \mathbf{x}^\star (note the \star instead of \sharp) as follows:

$$\mathbf{x}_J^\star \stackrel{\text{def}}{=} \mathbf{x}_J^\sharp - \lambda(\mathbf{A}_J^\top \mathbf{A}_J)^{-1} \text{sign}(\mathbf{x}_J^\sharp), \quad \mathbf{x}_{J^c}^\star \stackrel{\text{def}}{=} \mathbf{0} \quad (11)$$

There is a critical value $\lambda_J > 0$ such that for $\lambda < \lambda_J$ we have $\text{sign}(\mathbf{x}_J^\star) = \text{sign}(\mathbf{x}_J^\sharp)$, so that

$$\mathbf{A}_J^\top (\mathbf{y} - \mathbf{A}\mathbf{x}^\star) = \mathbf{A}_J^\top (\mathbf{A}_J \mathbf{x}_J^\sharp - \mathbf{A}_J \mathbf{x}_J^\star) = \lambda \mathbf{A}_J^\top \mathbf{A}_J (\mathbf{A}_J^\top \mathbf{A}_J)^{-1} \text{sign}(\mathbf{x}_J^\sharp) = \lambda \text{sign}(\mathbf{x}_J^\sharp) = \lambda \text{sign}(\mathbf{x}_J^\star).$$

This shows that $J \subseteq I(\mathbf{y}, \mathbf{A}, \lambda)$. Since \mathbf{A} is in general position with respect to \mathbf{y} and $\mathbf{y} \in \mathcal{R}(\mathbf{A}_J)$, we have $|J| \geq m$. We also know that, since \mathbf{A} has columns in general position, $|I| \leq m$ by (9). Hence $I = J$. This shows that \mathcal{J} is made of a *single* support set J , which matches the equicorrelation set $I(\mathbf{y}, \mathbf{A}, \lambda)$ for any small enough λ . As a consequence, \mathcal{X} is a singleton and (3) has a unique solution. \square

2.2 Generic Uniqueness and Optimal Sparsity of Spin

In the previous subsection we proved two facts about (vector) ℓ^1 minimization: that the optimizer is generically sparse and that it is unique. We now combine these results with known facts about optimal sparsity of dual frames. Let us first state a simple matrix consequence of Lemma 2.1:

Lemma 2.3. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$, $m < n$, and $0 \leq p \leq 1$. Then there exists $\mathbf{X} \in \text{ginv}_p(\mathbf{A})$ with at most m^2 nonzero entries (more precisely, with at most m nonzero elements per column). In particular, there is such an $\mathbf{X} \in \text{spin}(\mathbf{A}) = \text{ginv}_1(\mathbf{A})$.*

Remark 2.1. *One consequence of Lemma 2.3 is that $\text{spin}(\mathbf{A})$ is in general not the MPP \mathbf{A}^\dagger .*

Proof. Minimization for $\text{ginv}_p(\mathbf{A})$ can be decoupled into ℓ^p minimizations for every column,

$$\text{ginv}_p(\mathbf{A})_j = \arg \min_{\mathbf{Ax}=\mathbf{e}_j} \|\mathbf{x}\|_p. \quad (12)$$

By Lemma 2.1, the set of minimizers for every column of $\text{ginv}_p(\mathbf{A})$ (the set of solutions to (12)) contains a point with at most m nonzeros. Because $\text{ginv}_p(\mathbf{A})$ has m columns, there exists an element of $\text{ginv}_p(\mathbf{A})$ with at most m^2 non-zero entries. \square

A corollary of Lemma 2.3 is that generically, $\text{spin}(\mathbf{A})$ is *the* sparsest pseudoinverse of \mathbf{A} . To see this, we invoke a result by Krahmer, Kutyniok and Lemvig [KKL12]:

Lemma 2.4 ([KKL12, Theorem 3.6]). *Let $\mathcal{F}(m, n)$ be the set of full rank matrices in $\mathbb{R}^{m \times n}$, and denote by $\mathcal{N}(m, n)$ the subset of matrices in $\mathcal{F}(m, n)$ whose sparsest generalized inverse has m^2 non-zeros. Then*

- (i) *Any matrix in $\mathcal{F}(m, n)$ is arbitrarily close to a matrix in $\mathcal{N}(m, n)$,*
- (ii) *The set $\mathcal{F}(m, n) \setminus \mathcal{N}(m, n)$ has measure zero.*

In particular, the sparsest generalized inverse of many random matrices (e.g., iid Gaussian) will have m^2 non-zeros with probability 1.

Lemmata 2.3 and 2.4 imply that for almost all matrices the set of sparsest generalized inverses intersects the set of ℓ^1 -norm minimizing generalized inverses:

Theorem 2.2. *The set of matrices \mathbf{A} in $\mathcal{F}(m, n)$ such that $\text{spinv}(\mathbf{A})$ does not intersect $\text{ginv}_0(\mathbf{A})$ has measure zero.*

With all these results in hand we can prove a stronger claim which is our main result. Namely, the spinv is generically unique and m^2 -sparse.

Theorem 2.3. *Assume that $\mathbf{A} \in \mathbb{R}^{m \times n}$ has columns in general position, and that \mathbf{A} is in general position with respect to the canonical basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_m$. Then the sparse pseudoinverse $\text{spinv}(\mathbf{A})$ of \mathbf{A} contains a single matrix whose columns are all exactly m -sparse.*

Corollary 2.1. *Except for a measure-zero set of matrices, the sparse pseudoinverse $\text{spinv}(\mathbf{A})$ of $\mathbf{A} \in \mathbb{R}^{m \times n}$ contains a single matrix whose columns are all exactly m -sparse.*

Proof of Theorem 2.3. To finish the proof of Theorem 2.3, we apply Lemma 2.1 and Lemma 2.2 to the sparse pseudoinverse by setting $\mathbf{y} = \mathbf{e}_i$ to obtain \mathbf{x}_i , $i \leq 1 \leq m$. \square

Proof of Corollary 2.1. When \mathbf{A} is a random matrix drawn from a probability distribution which is absolutely continuous with respect to the Lebesgue measure, we have with probability one that: a) its columns are in general position; b) it is in general position with respect to any finite collection of vectors $\{\mathbf{y}_1, \dots, \mathbf{y}_\ell\}$. We can thus apply Theorem 2.3. \square

3 Numerical Stability

For matrices in general position (that is, for most matrices), there is a simpler way to obtain a generalized inverse with the minimal number m^2 of non-zeros than solving for spinv —just invert any full-rank $m \times m$ submatrix of \mathbf{A} . But there is a good reason why minimizing the ℓ^1 norm is a better idea than inverting a submatrix: it gives much better conditioned matrices.

To understand precisely what we mean by better conditioned, consider an overdetermined inverse problem $\mathbf{y} = \mathbf{B}\mathbf{x} + \mathbf{z}$, where $\mathbf{B} \in \mathbb{C}^{n \times m}$ is full rank with $m < n$, and \mathbf{z} is white noise. For any matrix $\mathbf{W} \in \mathcal{G}(\mathbf{B})$ we have that

$$\mathbb{E}_{\mathbf{z}}[\|\mathbf{W}\mathbf{y} - \mathbf{x}\|_2^2] = \mathbb{E}_{\mathbf{z}}[\|\mathbf{W}\mathbf{z}\|_2^2] \propto \|\mathbf{W}\|_F^2. \quad (13)$$

Thus the influence of noise on the output is controlled by the Frobenius norm of \mathbf{W} which makes it clear that it is desirable to use generalized inverses with small Frobenius norms. The right figure of merit is how much larger the Frobenius norm of our generalized inverse is than that of the MPP which attains the smallest one.

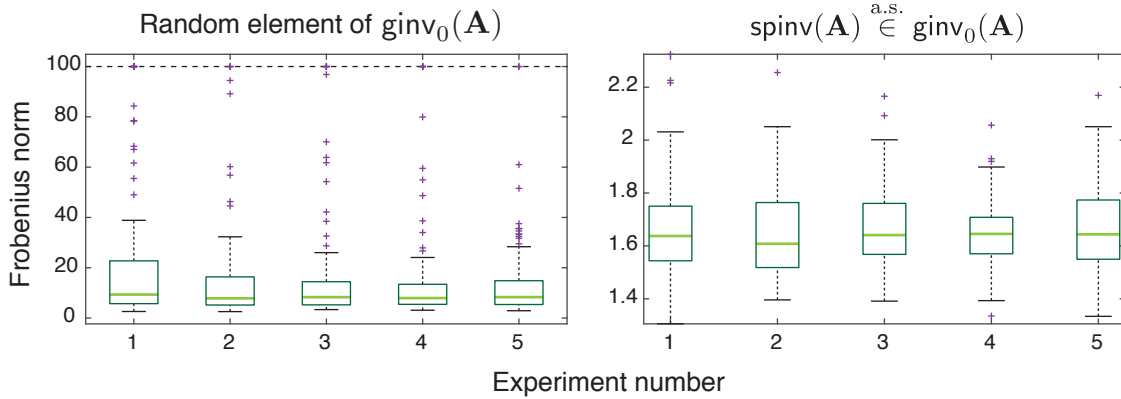


Figure 2: Frobenius norm of a random element of $\text{ginv}_0(\mathbf{A})$ for \mathbf{A} a random Gaussian matrix of size 20×30 and a sparse pseudoinverse $\text{spinv}(\mathbf{A})$ (generically in $\text{ginv}_0(\mathbf{A})$; cf. Corollary 2.1) of the same matrix. The random element of $\text{ginv}_0(\mathbf{A})$ is computed by selecting m columns of \mathbf{A} at random and inverting the obtained submatrix. The plot shows results of 5 identical experiments, each consisting of generating a random 20×30 matrix and computing the two inverses 100 times. In the first experiment, the outlier norms extend to 2000 so they were clipped at 100. Green lines denote medians, and boxes denote the second and the third quartile.

Alas, Figure 2 shows that for a simple inversion of an $m \times m$ submatrix, this Frobenius norm can be quite large. Worse, as figure shows, the variance of the norm of this inverted minor is large (note that the ordinate axis was clipped). On the contrary, if we carefully select a particular element of $\text{ginv}_0(\mathbf{A})$, namely the sparse pseudoinverse, then it appears that we get a well-controlled Frobenius norm; in other words we get a well-conditioned generalized inverse. The goal of this section is to make the last statement rigorous by developing concentration results for the Frobenius norm of $\text{spinv}(\mathbf{A})$ when \mathbf{A} is an iid Gaussian random matrix. In particular, we will prove the following result:

Theorem 3.1. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be a standard iid Gaussian matrix, $1 \leq m \leq n$. Define $\delta \stackrel{\text{def}}{=} (m-1)/n \in (0, 1)$ and for $1 \leq p \leq 2$ define the function*

$$D(t) = D_p(t; n) \stackrel{\text{def}}{=} \frac{1}{n} \left[\mathbb{E}_{\mathbf{h} \sim \mathcal{N}(0, \mathbf{I}_n)} \text{dist}(\mathbf{h}, \|\cdot\|_{p^*} \leq t) \right]^2 \in (0, 1), \quad (14)$$

Let $t^* = t_p^*(\delta, n)$ be the unique solution of $\delta = D(t) - \frac{t}{2}D'(t)$ on $(0, \infty)$ and denote

$$\alpha^* = \alpha_p^*(\delta, n) \stackrel{\text{def}}{=} \sqrt{\frac{D(t^*)}{\delta(\delta - D(t^*))}}.$$

If there exist $\gamma(\delta)$ and $N(\delta)$ such that $-t^*D'(t^*; n) \geq \gamma(\delta) > 0$ for all $n \geq N(\delta)$, then for any $n \geq \max(2/(1-\delta), N(\delta))$ we have: for any $0 < \epsilon \leq 1$

$$\mathbb{P} \left[\left| \frac{n}{m} \|\text{ginv}_p(\mathbf{A})\|_F^2 - (\alpha^*)^2 \right| \geq \epsilon(\alpha^*)^2 \right] \leq \frac{n}{C_1 \epsilon} e^{-C_2 n \epsilon^4}, \quad (15)$$

where the constants $C_1, C_2 > 0$ may depend on δ but not on n or ϵ .

The statement of Theorem 3.1 may be difficult to parse at first glance. Luckily, it allows us to obtain more explicit results for the two most interesting cases: $p = 1$ and $p = 2$. For $p = 2$ we get a result on the Frobenius norm of the Moore-Penrose pseudoinverse. It is complementary to a known large deviation bound [HMT11, Proposition A.5; Theorem A.6] obtained by a completely different technique.

Corollary 3.1 ($p = 2$). *With the notations of the above theorem, for $0 < \delta < 1$ and $n \geq N(\delta)$ we have: for any $0 < \epsilon \leq 1$*

$$\mathbb{P} \left[\left| \frac{n}{m} \|\mathbf{A}^\dagger\|_F^2 - (\alpha^*)^2 \right| \geq \epsilon (\alpha^*)^2 \right] \leq \frac{n}{C_1 \epsilon} e^{-C_2 n \epsilon^4}, \quad (16)$$

where the constants $C_1, C_2 > 0$ may depend on δ but not on n or ϵ , and $\alpha^* = \alpha_2^*(\delta; n)$ with

$$\lim_{n \rightarrow \infty} \alpha_2^*(\delta; n) = \frac{1}{\sqrt{1-\delta}}. \quad (17)$$

Remark 3.1. *This corollary covers “small” deviations ($0 < \epsilon \leq 1$). In contrast, the result of [HMT11] establishes that $\mathbb{E} \frac{n}{m} \|\mathbf{A}^\dagger\|_F^2 = \frac{n}{n-m-1} = \frac{1}{1-\delta}$ and that for any $\tau \geq 1$,*

$$\mathbb{P} \left[\frac{n}{m} \|\mathbf{A}^\dagger\|_F^2 \geq \frac{12n}{n-m} \tau \right] \leq 4\tau^{-(n-m)/4} = 4e^{-n \frac{(1-m/n) \log \tau}{4}}.$$

For large n we have $\frac{12n}{n-m} \tau \approx \frac{12\tau}{1-\delta} \approx 12\tau(\alpha_2^*)^2$ and $\frac{(1-m/n) \log \tau}{4} \approx \frac{(1-\delta) \log \tau}{4}$, hence this provides a bound for $1 + \epsilon := 12\tau \geq 12$ with an exponent $\log \tau$, of the order of $\log \epsilon$ (instead of ϵ^4 we get for $0 < \epsilon \leq 1$). Furthermore, unlike [HMT11], we give a two-sided bound; that is, we also show that the probability of $\frac{n}{m} \|\mathbf{A}^\dagger\|_F^2$ being much smaller than $(\alpha^*)^2$ is exponentially small.

The most interesting corollary is for $p = 1$.

Corollary 3.2 ($p = 1$). *With the notations of the above theorem, for $0 < \delta < 1$ and $n \geq N(\delta)$ we have: for any $0 < \epsilon \leq 1$*

$$\mathbb{P} \left[\left| \frac{n}{m} \|\text{spinv}(\mathbf{A})\|_F^2 - (\alpha^*)^2 \right| \geq \epsilon (\alpha^*)^2 \right] \leq \frac{n}{C_1 \epsilon} e^{-C_2 n \epsilon^4}, \quad (18)$$

where the constants $C_1, C_2 > 0$ may depend on δ but not on n or ϵ , and $\alpha^* = \alpha_1^*(\delta; n)$ with

$$\lim_{n \rightarrow \infty} \alpha_1^*(\delta; n) = \sqrt{\frac{1}{\sqrt{\frac{2}{\pi}} e^{-\frac{(t_1^*)^2}{2}} t_1^* - \delta (t_1^*)^2} - \frac{1}{\delta}} \quad (19)$$

$$t_1^* = \sqrt{2} \cdot \text{erfc}^{-1}(\delta).$$

Results of Corollaries 3.1 and 3.2 are illustrated in Figures 3 and 4. Figure 3 shows the shape of the limiting $(\alpha_p^*)^2$ as a function of δ , as well as empirical averages of 10 realization for different values of n and δ . As expected, the limiting values get closer to the empirical result as n grows larger; for $n = 1000$ the agreement is near-perfect.

In Figure 3 we also show the individual realizations for different combinations of n and δ . As predicted by the two corollaries, the variance of both $\frac{n}{m} \|\mathbf{A}^\dagger\|_F^2$ and $\frac{n}{m} \|\text{spinv}(\mathbf{A})\|_F^2$ reduces with n . For larger values of n , all realizations are very close to the limiting $(\alpha_p^*)^2$.

It is worth noting that the `spinv` and the MPP exhibit rather different qualitative behaviors. The Frobenius norm of the MPP monotonically decreases as δ gets smaller, while that of the `spinv` turns up below some critical δ . This may be understood by noting that for small δ , the support of the `spinv` is concentrated on very few entries, so those have to be comparably large to produce the ones in the identity matrix $\mathbf{I} = \mathbf{A}\mathbf{X}$. However, the ℓ^2 norm is more punishing for large entries.




Figure 3: Comparison of the limiting $(\alpha_p^*)^2$ with the mean of $\frac{n}{m} \|\mathbf{A}^\dagger\|_F^2$ and $\frac{n}{m} \|\text{spinv}(\mathbf{A})^\dagger\|_F^2$ for 10 realizations of \mathbf{A} . Empirical results are given for $\delta \in \{0.1, 0.2, \dots, 0.9\}$ and $n \in \{100, 200, 500, 1000\}$. Black squares represent the empirical result for $n = 1000$; colored squares represent the empirical mean of $\frac{n}{m} \|\mathbf{A}^\dagger\|_F^2$ for $n \in \{100, 200, 500\}$, with the largest discrepancy (purple squares) for $n = 100$; colored diamonds represent the empirical mean of $\frac{n}{m} \|\text{spinv}(\mathbf{A})^\dagger\|_F^2$ with the largest discrepancy (orange diamonds) again for $n = 100$.

Several remarks are due:

1. The bound (15) and the corresponding bounds in Corollaries 3.2 and 3.1 involve ϵ^4 instead of the usual ϵ^2 implying a higher variance of $\frac{n}{m} \|\text{ginv}_p(\mathbf{A})\|_F^2$: to guarantee a given probability in the right hand side of (15), ϵ should be of the order $n^{-1/4}$ instead of the usual $n^{-1/2}$. This seems to be a consequence of the technique used to lower bound $\inf_{|\alpha - \alpha^*| \geq \epsilon} \kappa(\alpha) - \kappa(\alpha^*)$ in Lemma B.7 which relies on strong convexity of κ . A more refined bound may give the result with better error bars, albeit also complicate the analysis.
2. Corollaries 3.1 and 3.2 specialize Theorem 3.1 to cases $p = 1$ and $p = 2$ proving that the Frobenius norm of the corresponding generalized inverses indeed concentrates, and giving a closed-form limiting value of the optimal α^* . It would seem rather natural that an interpolation to $p \in (1, 2)$ is possible, although α_p^* would be specified implicitly and have to be computed numerically. Whether an extension to $p > 2$ is possible is less clear. Numerical evidence suggests that it is, but we leave the corresponding theory to future work.




Figure 4: Comparison of the limiting $(\alpha_p^*)^2$ with the value of $\frac{n}{m} \|\text{spinv}(\mathbf{A})\|_F^2$ (left) and $\frac{n}{m} \|\mathbf{A}^\dagger\|_F^2$ (right) for 10 realizations of \mathbf{A} . Results are shown for $\delta \in \{0.1, 0.2, 0.4, 0.7\}$ and different values of n . Values for individual realizations are shown with colored dots with different color for every combination of n and δ . Horizontal dashed lines indicate the limiting value for the considered values of δ .

3. With additional work, one could characterize the rate of convergence of $\alpha_p^*(\delta; n)$ towards its limit in Corollaries 3.1 and 3.2. We leave these characterizations to future work.

Proof of Corollaries 3.1-3.2. For $p \in \{1, 2\}$, using Lemmas B.5-B.6 we lower bound $-t^* D'_p(t^*) \geq \gamma(\delta)$ for all n above some $N(\delta)$, and control $\lim_{n \rightarrow \infty} \alpha_p^*(\delta; n)$. Applying Theorem 3.1 yields the conclusion. \square

3.1 Proof of the Main Concentration Result, Theorem 3.1

We prove Theorem 3.1 by first noting that the matrix optimization for $\text{ginv}_p(\mathbf{A})$ decouples into m vector optimizations, and then using the following vector result for each column of $\text{ginv}_p(\mathbf{A})$:

Lemma 3.1. *With notations and assumptions as in Theorem 3.1, we have for $0 < \epsilon' \leq 1$ and $n \geq \max(2/(1 - \delta), N(\delta))$*

$$\mathbb{P} [|\sqrt{n} \|\mathbf{x}^*\| - \alpha^*| \geq \epsilon' \alpha^*] \leq \frac{1}{K_1 \epsilon'} e^{-K_2 n \epsilon'^4}.$$

where $K_1, K_2 > 0$ may depend on δ but not on n or ϵ' .

The proof of this result is given in Section 3.2. It is based on the convex Gaussian min-max theorem [TAH16], cf Appendix A.3. However, because the squared Frobenius norm is a sum of m squared ℓ^2 column norms, an “in probability” result on vector norms would not suffice. To address this shortcoming we developed a number of technical lemmas that lead to a stronger concentration result. With Lemma 3.1 we now continue working to prove Theorem 3.1.

By definition, we have

$$\text{ginv}_p(\mathbf{A}) = \arg \min_{\mathbf{A}\mathbf{X}=\mathbf{I}} \|\text{vec}(\mathbf{X})\|_p = \arg \min_{\mathbf{A}\mathbf{X}=\mathbf{I}} \|\text{vec}(\mathbf{X})\|_p^p,$$

where we assume that the solution is unique (for $p = 1$ this is true with probability 1, so

we can condition on this event). This optimization decouples over columns of \mathbf{X} : denoting $\mathbf{X}^* = \text{ginv}_p(\mathbf{A})$ we have for the i th column that $\mathbf{x}_i^* = \arg \min_{\mathbf{Ax}=\mathbf{e}_i} \|\mathbf{x}\|_p$.

Lemma 3.1 tells us that $\sqrt{n} \|\mathbf{x}_i^*\|_2$ remains close to (α^*) . However, to exploit the additivity of the squared Frobenius norms over columns a more useful statement would be that $n \|\mathbf{x}_i^*\|_2^2$ remains close to $(\alpha^*)^2$. To show that this is indeed the case, we note that $(n \|\mathbf{x}^*\|^2 - (\alpha^*)^2) = (\sqrt{n} \|\mathbf{x}^*\| - \alpha^*)(\sqrt{n} \|\mathbf{x}^*\| + \alpha^*)$ so we can write for any $b > 0$

$$\mathbb{P} \left[\left| n \|\mathbf{x}^*\|^2 - (\alpha^*)^2 \right| \geq \epsilon' (\alpha^*)^2 \right] \leq \mathbb{P} \left[\left| \sqrt{n} \|\mathbf{x}^*\| - \alpha^* \right| \geq \epsilon' (\alpha^*)^2 / b \right] + \mathbb{P} \left[(\sqrt{n} \|\mathbf{x}^*\| + \alpha^*) \geq b \right].$$

By taking $b = 3\alpha^*$, we bound the second term as

$$\mathbb{P} \left[(\sqrt{n} \|\mathbf{x}^*\| + \alpha^*) \geq b \right] \leq \mathbb{P} \left[(\sqrt{n} \|\mathbf{x}^*\| - \alpha^*) \geq 1\alpha^* \right] \leq \frac{1}{K_1} e^{-K_2 n}$$

using Lemma 3.1. To bound the first term we again use Lemma 3.1 to finally obtain

$$\mathbb{P} \left[\left| n \|\mathbf{x}^*\|^2 - (\alpha^*)^2 \right| \geq \epsilon (\alpha^*)^2 \right] \leq \frac{1}{K_1} e^{-K_2 n} + \frac{3}{K_1 \epsilon} e^{-K_2 n (\epsilon/3)^4} \leq \frac{1}{C_1 \epsilon} e^{-C_2 n \epsilon^4}$$

with an appropriate choice of C_1, C_2 .

This characterizes the squared ℓ^2 norm of one column of the MPP. The squared Frobenius norm is a sum of m such terms $\sum_{i=1}^m \|\mathbf{x}_i\|^2$ which are not independent. Our goal is to show that

$$\frac{n}{m} \|\mathbf{X}^*\|_F^2 = \frac{n}{m} \sum_{i=1}^m \|\mathbf{x}_i^*\|^2$$

stays close to $(\alpha^*)^2$ as well. We work as follows:

$$\begin{aligned} \mathbb{P} \left[\left| \frac{n}{m} \|\mathbf{X}^*\|_F^2 - (\alpha^*)^2 \right| > \epsilon (\alpha^*)^2 \right] &= \mathbb{P} \left[\left| \sum_{i=1}^m (n \|\mathbf{x}_i^*\|^2 - (\alpha^*)^2) \right| > m \epsilon (\alpha^*)^2 \right] \\ &\leq \mathbb{P} \left[\sum_{i=1}^m \left| n \|\mathbf{x}_i^*\|^2 - (\alpha^*)^2 \right| > m \epsilon (\alpha^*)^2 \right] = (*). \end{aligned}$$

Now observe that if the sum of m terms is to be larger than $m\epsilon$, then at least one term must be larger than ϵ , so we can continue writing:

$$\begin{aligned} (*) &\leq \mathbb{P} \left[\exists k \in \{1, \dots, m\} : \left| n \|\mathbf{x}_k^*\|^2 - (\alpha^*)^2 \right| > \epsilon (\alpha^*)^2 \right] = \mathbb{P} \left[\bigcup_{i=1}^m \left\{ \left| n \|\mathbf{x}_i^*\|^2 - (\alpha^*)^2 \right| > \epsilon (\alpha^*)^2 \right\} \right] \\ &\leq \sum_{i=1}^m \mathbb{P} \left[\left| n \|\mathbf{x}_i^*\|^2 - (\alpha^*)^2 \right| > \epsilon (\alpha^*)^2 \right] \\ &\leq m \frac{1}{C_1 \epsilon} e^{-C_2 n \epsilon^4} \leq n \frac{1}{C_1 \epsilon} e^{-C_2 n \epsilon^4}, \end{aligned} \tag{20}$$

which completes the proof.

3.2 Proof of the Main Vector Result, Lemma 3.1

In the Appendix we establish Lemma B.1³ which lets us rewrite the optimization

$$\mathbf{x}^* = \arg \min_{\mathbf{Ax}=\mathbf{e}_1} \|\mathbf{x}\|_p \quad (21)$$

in an unconstrained form as (note that the data fidelity term is *not* squared here):

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{e}_1\| + \lambda \|\mathbf{x}\|_p. \quad (22)$$

More precisely, since the ℓ^p norm is L -Lipschitz with respect to the ℓ^2 norm (with $L \stackrel{\text{def}}{=} n^{\max(1/p-1/2, 0)}$), Lemma B.1 tells us that if we choose $\lambda \leq \frac{\sqrt{n}-\sqrt{m}}{L}(1-\epsilon)$, minimizers \mathbf{x}^* and $\tilde{\mathbf{x}}$ of (21) and (22) coincide⁴ with probability at least $1 - 2e^{-\epsilon^2(\sqrt{n}-\sqrt{m})^2/2}$. Rewriting the ℓ^2 norm in a variational form, we obtain

$$\tilde{\mathbf{x}} = \arg \min_{\mathbf{x}} \max_{\mathbf{u}: \|\mathbf{u}\| \leq 1} \mathbf{u}^\top \mathbf{Ax} + \lambda \|\mathbf{x}\|_p - \mathbf{u}^\top \mathbf{e}_1. \quad (23)$$

The expression (23) is a sum of a bilinear term involving \mathbf{A} and a convex-concave function⁵ $\psi(\mathbf{x}, \mathbf{u}) = \lambda \|\mathbf{x}\|_p - \mathbf{u}^\top \mathbf{e}_1$. That is exactly the structure required by the convex Gaussian min-max theorem [TO15, Theorem 3] [TAH16, Theorem 6.1]. For readers' convenience, we reproduce the version of CGMT we use in Appendix A.3 as Theorem A.1. The only non-conforming detail is that \mathbf{x} is not constrained to be in a compact set. To address this, instead of (23), we analyze the following bounded modification:

$$\tilde{\mathbf{x}}_K = \arg \min_{\mathbf{x}: \|\mathbf{x}\| \leq K} \max_{\mathbf{u}: \|\mathbf{u}\| \leq 1} \mathbf{u}^\top \mathbf{Ax} + \lambda \|\mathbf{x}\|_p - \mathbf{u}^\top \mathbf{e}_1 \quad (24)$$

We will show in due time that $\tilde{\mathbf{x}}_K = \tilde{\mathbf{x}}$ with high probability.

By the CGMT, instead of analyzing the so-called principal optimization (24), we can analyze an auxiliary optimization

$$\hat{\mathbf{x}}_K = \arg \min_{\mathbf{x}: \|\mathbf{x}\| \leq K} \max_{\mathbf{u}: \|\mathbf{u}\| \leq 1} \|\mathbf{x}\| \mathbf{g}^\top \mathbf{u} - \|\mathbf{u}\| \mathbf{h}^\top \mathbf{x} + \lambda \|\mathbf{x}\|_p - \mathbf{u}^\top \mathbf{e}_1, \quad (25)$$

where $\|\cdot\|$ are 2-norms and $\mathbf{g} \in \mathbb{R}^m$ and $\mathbf{h} \in \mathbb{R}^n$ are iid standard Gaussian random vectors. Part 2 of the CGMT tells us that if the optimal value of the auxiliary optimization (25) concentrates (note that (25) is a random optimization program), the optimal value of (24) will concentrate around the same value. This lets us prove that if the norm $\|\hat{\mathbf{x}}_K\|$ of the optimizer of (25) concentrates, the norm $\|\tilde{\mathbf{x}}_K\|$ of the optimizer of (24) will also concentrate around the same value.⁶

We will now formally show how to use this property for our purpose by going through a series of steps to simplify (25). The analysis will be easier if the argument of the optimization is of order 1, so let \mathbf{z} be an appropriately scaled version of \mathbf{x} , $\mathbf{z} = \mathbf{x}\sqrt{n}$ (accordingly $A = K\sqrt{n}$). Using the variational characterization of the ℓ^p norm we can rewrite (25) as

$$\hat{\mathbf{z}}_A = \arg \min_{\mathbf{z}: \|\mathbf{z}\| \leq A} \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{u}: \|\mathbf{u}\| = \beta \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \frac{1}{\sqrt{n}} \|\mathbf{z}\| \mathbf{g}^\top \mathbf{u} - \frac{1}{\sqrt{n}} \|\mathbf{u}\| \mathbf{h}^\top \mathbf{z} - \mathbf{u}^\top \mathbf{e}_1 + \frac{\lambda}{\sqrt{n}} \mathbf{w}^\top \mathbf{z}. \quad (26)$$

³We note that Lemma B.1 is a simple restatement of [OTH13, Lemma 9.2] with dependence on ϵ made explicit.

⁴An analogous result does not hold for the squared lasso (except for $\lambda = 0^+$).

⁵Convex in the first argument, concave in the second one.

⁶The referenced CGMT contains a similar statement, albeit we need a different derivation to get exponential concentration.

Let $\hat{\mathbf{z}}_A$ be a (random) optimizer of (26), and $\check{\mathbf{z}}_A$ a (random) value of \mathbf{z} at the optimum of the following program with the same cost function but an altered order of minimization and maximization:

$$\max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\mathbf{z}: \|\mathbf{z}\| \leq A} \max_{\mathbf{u}: \|\mathbf{u}\| = \beta} \frac{1}{\sqrt{n}} \|\mathbf{z}\| \mathbf{g}^\top \mathbf{u} - \frac{1}{\sqrt{n}} \|\mathbf{u}\| \mathbf{h}^\top \mathbf{z} - \mathbf{u}^\top \mathbf{e}_1 + \frac{\lambda}{\sqrt{n}} \mathbf{w}^\top \mathbf{z}. \quad (27)$$

Even though after performing optimization over \mathbf{u} in (26) we do not get a convex-concave cost function, we will see that we can indeed reorder maximizations as in (27) for our purposes. In the following we focus on (27) since it is easier to analyze. We compute as follows:

$$(27) = \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\mathbf{z}: \|\mathbf{z}\| \leq A} \beta \left\| \frac{\|\mathbf{z}\| \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\beta}{\sqrt{n}} \mathbf{h}^\top \mathbf{z} + \frac{\lambda}{\sqrt{n}} \mathbf{w}^\top \mathbf{z} \quad (28a)$$

$$= \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\alpha: 0 \leq \alpha \leq A} \min_{\mathbf{z}: \|\mathbf{z}\| = \alpha} \beta \left\| \frac{\|\mathbf{z}\| \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\beta}{\sqrt{n}} \mathbf{h}^\top \mathbf{z} + \frac{\lambda}{\sqrt{n}} \mathbf{w}^\top \mathbf{z} \quad (28b)$$

$$= \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\alpha: 0 \leq \alpha \leq A} \beta \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\alpha}{\sqrt{n}} \|\beta \mathbf{h} - \lambda \mathbf{w}\|. \quad (28c)$$

The objective in the last line is convex in α and jointly concave in (β, \mathbf{w}) . Additionally, the constraint sets are all convex and bounded, so by [Sio58, Corollary 3.3] we can swap the order of min and max again. We thus continue writing:

$$= \min_{\alpha: 0 \leq \alpha \leq A} \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \beta \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\alpha}{\sqrt{n}} \|\beta \mathbf{h} - \lambda \mathbf{w}\| \quad (29a)$$

$$= \min_{\alpha: 0 \leq \alpha \leq A} \max_{\beta: 0 \leq \beta \leq 1} \left(\beta \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\alpha}{\sqrt{n}} \min_{\mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1} \|\beta \mathbf{h} - \lambda \mathbf{w}\| \right) \quad (29b)$$

$$= \min_{\alpha: 0 \leq \alpha \leq A} \max_{\beta: 0 \leq \beta \leq 1} \beta \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\alpha \beta}{\sqrt{n}} \text{dist}(\beta \mathbf{h}, \|\cdot\|_{p^*} \leq \lambda) \quad (29c)$$

$$= \min_{\alpha: 0 \leq \alpha \leq A} \max_{\beta: 0 \leq \beta \leq 1} \phi(\alpha, \beta; \mathbf{g}, \mathbf{h}) \quad (29d)$$

where

$$\phi(\alpha, \beta; \mathbf{g}, \mathbf{h}) \stackrel{\text{def}}{=} \beta \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \frac{\alpha}{\sqrt{n}} \text{dist}(\beta \mathbf{h}, \|\cdot\|_{p^*} \leq \lambda). \quad (30)$$

We have achieved two important feats: 1) we simplified a high-dimensional vector optimization (22) into an optimization over two scalars (29d), 2) one of these scalars, $\check{\alpha} = \|\check{\mathbf{z}}_A\|$ is almost giving us what we seek—the (scaled) Frobenius norm of \mathbf{X} . To put all pieces together, there now remains to formally prove that the min-max switches and the concentration results we mentioned actually hold.

Combining the ingredients. By Lemma B.2, $\phi(\alpha, \beta; \mathbf{g}, \mathbf{h})$ concentrates around some deterministic function $\kappa(\alpha, \beta) = \kappa_p(\alpha, \beta; n, \delta, \lambda)$ with $\delta = (m-1)/n$. By Lemma B.7-Property 3, the minimizer of this function is given by

$$\arg \min_{\alpha: 0 \leq \alpha \leq A} \max_{\beta: 0 \leq \beta \leq 1} \kappa(\alpha, \beta) = \alpha^* = \alpha_p^*(\delta; n) \stackrel{\text{def}}{=} \sqrt{\frac{D_p(t_p^*, n)}{\delta(\delta - D_p(t_p^*, n))}} \quad (31)$$

as soon as $A = K\sqrt{n} > \alpha^*$ and $\lambda \leq t^*$ (cf. Lemma B.3-Property 7 for the definition of $t^* = t_p^*(\delta, n)$). By Lemma B.9, for $0 < \epsilon \leq \max(\alpha^*, A - \alpha^*)$, the minimizer

$$\alpha_\phi^* \stackrel{\text{def}}{=} \arg \min_{0 \leq \alpha \leq A} \max_{0 \leq \beta \leq 1} \phi(\alpha, \beta; \mathbf{g}, \mathbf{h})$$

stays ϵ -close to α^* with high probability:

$$\mathbb{P}[|\alpha_\phi^* - \alpha^*| \geq \epsilon] \leq \zeta(n, \frac{\omega(\epsilon)}{2}),$$

where

$$\zeta(n, \xi) = \zeta(n, \xi; A, \delta) \stackrel{\text{def}}{=} \frac{c_1}{\xi} e^{-c_2 \xi^2 n \cdot c(\delta, A)}, \quad \text{with } c(\delta, A) \stackrel{\text{def}}{=} \frac{\min(\delta, A^{-2})}{1 + \delta A^2}$$

and c_1, c_2 universal constants, and

$$\omega(\epsilon) = \omega_p(\epsilon; n, \delta, \lambda) \stackrel{\text{def}}{=} \frac{\epsilon^2}{2} \frac{\lambda \delta / t^*}{(1 + \delta(\alpha^* + \epsilon)^2)^{3/2}}.$$

As a consequence, as established in Lemma B.11, the scaled norm $\sqrt{n} \|\tilde{\mathbf{x}}_K\|$ of the minimizer of the (bounded) principal optimization problem (24) stays ϵ -close to α^* with high probability

$$\mathbb{P}[|\sqrt{n} \|\tilde{\mathbf{x}}_K\| - \alpha^*| \geq \epsilon] \leq 4\zeta(n, \frac{\omega(\epsilon)}{2}).$$

Similarly, for $0 < \epsilon \leq \min(\alpha^*, A - \alpha^*)$ by invoking Lemma B.12 we obtain that the scaled norm $\sqrt{n} \|\tilde{\mathbf{x}}\|$ of the minimizer of the *unbounded* optimization (23) stays in the neighborhood of α^* with high probability

$$\mathbb{P}[|\sqrt{n} \|\tilde{\mathbf{x}}\| - \alpha^*| \geq \epsilon] \leq 4\zeta(n, \frac{\omega(\epsilon)}{2})$$

and in fact $\mathbb{P}[\tilde{\mathbf{x}} \neq \tilde{\mathbf{x}}_K] \leq 4\zeta(n, \frac{\omega(\epsilon)}{2})$.

Since the ℓ^p norm is L -Lipschitz with respect to the Euclidean metric in \mathbb{R}^n , with $L = n^{\max(1/p-1/2, 0)}$, invoking Lemma B.1 yields that, for any

$$\lambda < \lambda_{\max}(n, m, t^*) \stackrel{\text{def}}{=} \min \left\{ \frac{\sqrt{n} - \sqrt{m}}{2L}, t^* \right\}.$$

the minimizer \mathbf{x}^* of the equality-constrained optimization (21) coincides with the minimizer $\tilde{\mathbf{x}}$ of the lasso formulation (22)-(23) except with probability at most $e^{-n(1-\sqrt{\delta+1/n})^2/8}$.

Overall this yields, for $\lambda < \lambda_{\max}$ and $\epsilon \leq \min(\alpha^*, A - \alpha^*)$:

$$\mathbb{P}[|\sqrt{n} \|\mathbf{x}^*\| - \alpha^*| \geq \epsilon] \leq e^{-n(1-\sqrt{\delta+1/n})^2/8} + 4\zeta(n, \frac{\omega(\epsilon)}{2}). \quad (32)$$

The infimum over admissible values of λ is obtained by taking its value when $\lambda = \lambda_{\max}$.

Expliciting the bound (32). From now on we choose $A \stackrel{\text{def}}{=} 2\alpha^*$ and, for $0 < \epsilon' \leq 1$, we consider $\epsilon \stackrel{\text{def}}{=} \epsilon' \alpha^*$ (which satisfies $0 < \epsilon \leq \min(\alpha^*, A - \alpha^*) = \alpha^*$). We use \lesssim_δ and \gtrsim_δ to denote inequalities up to a constant that may depend on δ , but not on n or ϵ' , provided $n \geq n(\delta)$. We specify $n(\delta)$ where appropriate.

By Lemma B.4-item 1, for any $1 \leq p \leq 2$ and any $n \geq 1$

$$\lambda_{\max} = t^* \min \left(\frac{\sqrt{n} - \sqrt{m}}{n^{1/p-1/2}} \frac{1}{2t^*}, 1 \right) \gtrsim_\delta t^* \min \left(\frac{\sqrt{n}(1-\sqrt{m/n})}{n^{1/p-1/2}} \frac{1}{n^{1-1/p}}, 1 \right) = t^* \min \left(1 - \sqrt{m/n}, 1 \right).$$

For $n \geq \frac{2}{1-\delta}$ we have

$$1 - \sqrt{m/n} = 1 - \sqrt{\delta + \frac{1}{n}} \geq 1 - \sqrt{(1+\delta)/2} \gtrsim_{\delta} 1 \quad (33)$$

hence $\lambda_{\max} \gtrsim_{\delta} t^*$.

With the shorthands $D(t) = D_p(t; n)$ and $D = D_p(t^*; n)$, we have

$$1 + \delta(\alpha^* + \epsilon)^2 \leq 1 + \delta A^2 \leq 4(1 + \delta(\alpha^*)^2) = 4(1 + \frac{D}{\delta-D}) = \frac{4\delta}{\delta-D} \lesssim_{\delta} (\delta - D)^{-1}.$$

hence with $\lambda = \lambda_{\max}$ we get for $n \geq 2/(1-\delta)$

$$\omega(\epsilon) = \frac{\epsilon'^2 \lambda_{\max}}{2 t^*} \frac{\delta(\alpha^*)^2}{(1 + \delta(\alpha^* + \epsilon)^2)^{3/2}} \gtrsim_{\delta} \epsilon'^2 \frac{\frac{D}{\delta-D}}{(\frac{4\delta}{\delta-D})^{3/2}} \gtrsim_{\delta} \epsilon'^2 D \sqrt{\delta - D}$$

By Lemma B.4-item 4 we have $D = D_p(t^*; n) \geq (\delta/C)^2$ for a universal constant C independent of n or p , hence for $n \geq 1$, $D \gtrsim_{\delta} 1$ and for $n \geq 2/(1-\delta)$,

$$\omega(\epsilon) \gtrsim_{\delta} \sqrt{-t^* D'_p(t^*; n)} \epsilon'^2$$

Moreover since $\min(\delta, A^{-2}) \geq \frac{1}{4} \min(\delta, (\alpha^*)^{-2}) = \frac{1}{4} \min(\delta, \delta(\delta - D)/D)$ we also get

$$c(\delta, A) = \frac{\min(\delta, A^{-2})}{1 + \delta A^2} \geq \frac{\min(\delta, \delta(\delta - D)/D)}{\delta/(\delta - D)} = \min(\delta - D, \frac{(\delta - D)^2}{D}) \geq (\delta - D)^2 = \frac{1}{4} [-t^* D'_p(t^*; n)]^2.$$

Since $-t^* D'_p(t^*; n) \geq \gamma(\delta) > 0$ for any $n \geq N(\delta)$ (recall that $t^* = t_p^*(\delta; n)$), we have

$$-t^* D'_p(t^*) \gtrsim_{\delta} 1, \quad (34)$$

and we obtain for $n \geq \max(2/(1-\delta), N)$: $\omega(\epsilon) \gtrsim_{\delta} \epsilon'^2$ and $c(\delta, A) \gtrsim_{\delta} 1$. Combining the above yields, for $0 < \epsilon' \leq 1$, $n \geq \max(2/(1-\delta), N(\delta))$:

$$\begin{aligned} \mathbb{P}\left[\left|\|\mathbf{x}^*\| - \frac{\alpha^*}{\sqrt{n}}\right| \geq \frac{\epsilon' \alpha^*}{\sqrt{n}}\right] &\leq e^{-C_1 n} + 4\zeta(n, C_2 \epsilon'^2) \\ &\leq \frac{1}{K_1(\epsilon')^2} e^{-K_2(\epsilon')^4 n} \end{aligned}$$

with $K_i, C_i \gtrsim_{\delta} 1$.

4 Conclusion

In this paper (Part II of the ‘‘Beyond Moore-Penrose’’ mini-series) we looked at generalized matrix inverses which minimize entrywise ℓ^p norms, with a particular emphasis on the sparse pseudoinverse ($p = 1$). Our central result is Theorem 3.1 together with Corollaries 3.1 and 3.2 which discuss numerical stability of a class of generalized inverses as measured by their Frobenius norm. This allows us to quantify the MSE hit incurred by using the sparse pseudoinverse instead of the MPP and to show that, in fact, this hit is controlled.

We highlight three main results about the sparse pseudoinverse $\text{spinv}(\mathbf{A}) = \text{ginv}_1(\mathbf{A})$ of a *generic* matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$: 1) it is unique, 2) it has precisely m zeros per column, and 2) its Frobenius norm is characterized by Corollary 3.2. For a large range of m/n (with $m < n$) the Frobenius norm of the sparse pseudoinverse is relatively close to the Frobenius norm of the MPP

which is the smallest possible among all generalized inverses (Figure 3). This does not hold for certain ad hoc strategies that yield generalized inverses with the same non-zero count (Figure 2). More generally, we gave finite-size concentration bounds for the square of the Frobenius norm of all $\text{ginv}_p(\cdot)$, $1 \leq p \leq 2$, complementing a known result for $p = 2$.

Theorem 3.1 of Part I and Corollary 3.2 together establish a form of regularity of the sparse pseudoinverse, namely that $\frac{n}{m}\mathbb{E}[\text{spinv}(\mathbf{A})\mathbf{A}\mathbf{x}] = \mathbf{x}$ and that $\|\text{spinv}(\mathbf{A})\mathbf{A}\mathbf{x}\|$ is well-behaved (at least for Gaussian random matrices)—two essential properties exhibited by the MPP.

As a useful side product, along the way we presented a number of new results related to matrix analysis, sparse representations, and convex programming. For example, we proved that basis pursuit generically has a unique minimizer—a folklore fact mentioned repeatedly in the literature, but for which we could find no proof. The same goes for the fact that ℓ_p minimization for $0 \leq p \leq 1$ has a sparse solution—to the best of our knowledge, yet another piece of thus far unproven folklore.

The most important future work is to develop extensions of Theorem 3.1 and the two corollaries to $p > 2$.

5 Acknowledgments

The authors would like to thank Mihailo Kolundžija, Miki Elad, Jakob Lemvig, and Martin Vetterli for the discussions and input in preparing this manuscript. A special thanks goes to Christos Thrampoulidis for his help in understanding and applying the Gaussian min-max theorem, and to Simon Foucart for discussions leading to the proof of Lemma 2.1 for $p < 1$.

This work was supported in part by the European Research Council, PLEASE project (ERC-StG-2011-277906).

Appendices

A Results about Gaussian processes

A.1 Concentration of measure

Lemma A.1. *Let \mathbf{h} be a standard Gaussian random vector of length n , $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ a 1-Lipschitz function. Then the following hold:*

(a) *For any $0 < \epsilon < 1$, $\mathbb{P}[\|\mathbf{h}\|^2 \leq n(1 - \epsilon)] \leq e^{-\epsilon^2/4}$; $\mathbb{P}[\|\mathbf{h}\|^2 \geq n/(1 - \epsilon)] \leq e^{-\epsilon^2/4}$;*

(b) *For any $0 < \epsilon < 1$, $\mathbb{P}\left\{\|\mathbf{h}\|^2 \notin [(1 - \epsilon)n, n/(1 - \epsilon)]\right\} \leq 2e^{-\frac{\epsilon^2 n}{4}}$;*

(c) *For any $\epsilon > 0$, $\mathbb{P}\left[\left|\|\mathbf{h}\|^2 - n\right| \geq \sqrt{\epsilon n}\right] \leq \begin{cases} 2e^{-\frac{\epsilon}{8}} & \text{for } 0 \leq \epsilon \leq n, \\ 2e^{-\frac{\sqrt{\epsilon n}}{8}} & \text{for } \epsilon > n. \end{cases}$*

(d) *For any $u > 0$, $\mathbb{P}[f(\mathbf{h}) - \mathbb{E}f(\mathbf{h}) \geq u] \leq e^{-u^2/2}$; $\mathbb{P}[f(\mathbf{h}) - \mathbb{E}f(\mathbf{h}) \leq -u] \leq e^{-u^2/2}$;*

(e) *For any $u > 0$, $\mathbb{P}[|f(\mathbf{h}) - \mathbb{E}f(\mathbf{h})| \geq u] \leq 2e^{-u^2/2}$;*

(f) $\text{Var}[f(\mathbf{h})] \leq 1$

Proofs and references. We mostly refer to texts where the proofs can be found and show how to modify the standard forms of the bounds address our needs.

(a) See [Bar05, Corollary 2.3].

(b) Immediate consequence of (a) by a union bound.

(c) We use a Bernstein-type inequality for sub-exponential random variables. Namely, $X_i = h_i^2$ is χ_1^2 which is subexponential with parameters $\nu = 2, b = 4$ [Wai15, Example 2.4], i.e. $\mathbb{E}[e^{\lambda(X_i - \mu)}] \leq e^{\frac{\nu^2 \lambda^2}{2}}$ for all $|\lambda| \leq \frac{1}{b}$, and $\mu = \mathbb{E}[X_i] = 1$. Applying [Wai15, Proposition 2.2] then yields

$$\mathbb{P} \left[\left| \|\mathbf{h}\|^2 - n \right| \geq \sqrt{\epsilon n} \right] \leq 2 \begin{cases} e^{-\frac{\epsilon}{8}} & \text{for } 0 \leq \epsilon \leq n, \\ e^{-\frac{\sqrt{\epsilon n}}{8}} & \text{for } \epsilon > n. \end{cases}$$

(d) [Led99, Eq. (1.22)].

(e) Union bound applied to (d).

(f) A consequence of Poincaré inequality for Gaussian measures [Led99, Eq. (2.16)]: $\text{Var}[f(\mathbf{h})] \leq \mathbb{E}[\|\nabla f(\mathbf{h})\|^2]$ for 1-Lipschitz f for which $\|\nabla f(\mathbf{h})\| \leq 1$.

□

A.2 Complementary error function and related functions

We will be using the complementary error function and will need some basic properties which can be found e.g. in [CDS03] or references therein.

Lemma A.2. *The complementary error function is defined for any $z \in \mathbb{R}$ as*

$$\text{erfc}(z) \stackrel{\text{def}}{=} \frac{2}{\sqrt{\pi}} \int_z^\infty e^{-t^2} dt. \quad (35)$$

For any $z > 0$ we have

$$\text{erfc}(z) \leq \exp(-z^2). \quad (36)$$

Moreover, if $g \sim \mathcal{N}(0, 1)$ is a standard centered normal variable then any $x > 0$

$$\mathbb{P} \{ |g| > x \} = \text{erfc}(x/\sqrt{2}). \quad (37)$$

Lemma A.3. *Let $\theta(t) \stackrel{\text{def}}{=} \mathbb{E}(|h| - t)_+^2$, where $h \sim \mathcal{N}(0, 1)$ and $t \geq 0$. Then*

1. $\theta(t) = (t^2 + 1) \text{erfc}\left(\frac{t}{\sqrt{2}}\right) - \sqrt{\frac{2}{\pi}} e^{-\frac{t^2}{2}} t$,
2. $\theta(t) - (t/2)\theta'(t) = \text{erfc}(t/\sqrt{2})$.

Proof. $\theta(t) = \mathbb{E}(|h| - t)_+^2 = \int_{-\infty}^\infty (|\varphi| - t)_+^2 p_h(\varphi) d\varphi = 2 \int_t^\infty (\varphi^2 - 2t\varphi + t^2) p_h(\varphi) d\varphi$.

Now we compute as follows:

$$\begin{aligned}
\bullet \int_t^\infty \varphi^2 p_h(\varphi) \, d\varphi &= \frac{1}{\sqrt{2\pi}} \int_t^\infty \varphi^2 e^{-\frac{\varphi^2}{2}} \, d\varphi = \frac{1}{\sqrt{2\pi}} \int_t^\infty (-\varphi) \frac{d}{d\varphi} \left\{ e^{-\frac{\varphi^2}{2}} \right\} \, d\varphi \\
&= \frac{1}{\sqrt{2\pi}} \left(\left[-\varphi e^{-\frac{\varphi^2}{2}} \right]_t^\infty + \int_t^\infty e^{-\frac{\varphi^2}{2}} \, d\varphi \right) \\
&\stackrel{(a)}{=} \frac{1}{2} \operatorname{erfc} \left(\frac{t}{\sqrt{2}} \right) + \frac{e^{-\frac{t^2}{2}}}{\sqrt{2\pi}},
\end{aligned}$$

where in (a) we used Lemma A.2,

$$\begin{aligned}
\bullet \int_t^\infty \varphi p_h(\varphi) \, d\varphi &= \frac{1}{\sqrt{2\pi}} \int_t^\infty \frac{d}{d\varphi} \left\{ -e^{-\frac{\varphi^2}{2}} \right\} \, d\varphi = \frac{1}{\sqrt{2\pi}} \left[-e^{-\frac{\varphi^2}{2}} \right]_t^\infty = \frac{e^{-\frac{t^2}{2}}}{\sqrt{2\pi}}, \\
\bullet \int_t^\infty p_h(\varphi) \, d\varphi &= \frac{1}{2} \operatorname{erfc} \left(\frac{t}{\sqrt{2}} \right) \text{ again by Lemma A.2.}
\end{aligned}$$

Putting together the above we get the desired expression. The second statement follows from the fact that $\frac{d}{dt} \operatorname{erfc} \left(\frac{t}{\sqrt{2}} \right) = -\sqrt{\frac{2}{\pi}} e^{-\frac{t^2}{2}}$. \square

A.3 Convex Gaussian Min-Max Theorem (CGMT) [TAH16, Theorem 6.1]

Let the principal optimization (PO) and auxiliary optimization (AO) be defined as

$$\Phi(\mathbf{G}) \stackrel{\text{def}}{=} \min_{\mathbf{v} \in \mathcal{S}_{\mathbf{v}}} \max_{\mathbf{u} \in \mathcal{S}_{\mathbf{u}}} \mathbf{u}^\top \mathbf{G} \mathbf{v} + \psi(\mathbf{v}, \mathbf{u}) \quad (38)$$

$$\phi(\mathbf{g}, \mathbf{h}) \stackrel{\text{def}}{=} \min_{\mathbf{v} \in \mathcal{S}_{\mathbf{v}}} \max_{\mathbf{u} \in \mathcal{S}_{\mathbf{u}}} \|\mathbf{v}\|_2 \mathbf{g}^\top \mathbf{u} + \|\mathbf{u}\|_2 \mathbf{h}^\top \mathbf{v} + \psi(\mathbf{v}, \mathbf{u}), \quad (39)$$

with $\mathbf{G} \in \mathbb{R}^{m \times n}$, $\mathbf{g} \in \mathbb{R}^m$, $\mathbf{h} \in \mathbb{R}^n$, $\mathcal{S}_{\mathbf{v}} \subset \mathbb{R}^n$, $\mathcal{S}_{\mathbf{u}} \subset \mathbb{R}^m$ and $\psi : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$. Let further $\mathbf{v}_\Phi \stackrel{\text{def}}{=} \mathbf{v}_\Phi(\mathbf{G})$ denote any optimal minimizer of (PO) and $\mathbf{v}_\phi \stackrel{\text{def}}{=} \mathbf{v}_\phi(\mathbf{g}, \mathbf{h})$ any optimal minimizer of (AO). Then the following holds:

Theorem A.1. *In (38) and (39), let $\mathcal{S}_{\mathbf{v}}$ and $\mathcal{S}_{\mathbf{u}}$ be compact and ψ continuous on $\mathcal{S}_{\mathbf{v}} \times \mathcal{S}_{\mathbf{u}}$. Let also \mathbf{G} , \mathbf{g} , \mathbf{h} have entries that are iid standard normal. The following hold:*

1. For all $c \in \mathbb{R}$

$$\mathbb{P}[\Phi(\mathbf{G}) < c] \leq 2\mathbb{P}[\phi(\mathbf{g}, \mathbf{h}) \leq c]$$

2. Assume further that $\psi(\mathbf{v}, \mathbf{u})$ is convex-concave on $\mathcal{S}_{\mathbf{v}} \times \mathcal{S}_{\mathbf{u}}$ where $\mathcal{S}_{\mathbf{v}}$ and $\mathcal{S}_{\mathbf{u}}$ are convex. Then for all $c \in \mathbb{R}$

$$\mathbb{P}[\Phi(\mathbf{G}) > c] \leq 2\mathbb{P}[\phi(\mathbf{g}, \mathbf{h}) \geq c].$$

In particular, for all $\mu \in \mathbb{R}$ and $t > 0$,

$$\mathbb{P}[|\Phi(\mathbf{G}) - \mu| > t] \leq 2\mathbb{P}[|\phi(\mathbf{g}, \mathbf{h}) - \mu| \geq t].$$

B Lemmata for Section 3

Lemma B.1 (Rewriting of [OTH13, Lemma 9.2] with dependence on ϵ made explicit). *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be a random matrix with iid standard normal entries, and $m < n$. Let further*

$\mathbf{y} \in \mathbb{R}^m$. Consider the solution of an ℓ^2 -lasso with a regularizer f which is L -Lipschitz with respect to the ℓ^2 -norm:

$$\mathbf{x}^* \stackrel{\text{def}}{=} \arg \min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{y} - \mathbf{A}\mathbf{x}\| + \lambda f(\mathbf{x}).$$

Then for any $0 \leq \epsilon < 1$ and $0 < \lambda < \frac{\sqrt{n} - \sqrt{m}}{L}(1 - \epsilon)$ the solution \mathbf{x}^* satisfies $\mathbf{y} = \mathbf{A}\mathbf{x}^*$ with probability at least $1 - e^{-\epsilon^2(\sqrt{n} - \sqrt{m})^2/2}$. In other words, ℓ^2 -lasso (let us emphasize that the data fidelity term is not squared here) gives the same result as the equality-constrained f -minimization.

Proof. Using [Ver09, Corollary 5.35]⁷, we have for every $t \geq 0$

$$\mathbb{P}[\sigma_{\min}(\mathbf{A}^\top) \leq \sqrt{n} - \sqrt{m} - t] \leq e^{-t^2/2}.$$

Dividing by $\sqrt{n} - \sqrt{m}$, we can rewrite this as

$$\mathbb{P}[\sigma_{\min}(\mathbf{A}^\top)/(\sqrt{n} - \sqrt{m}) \leq 1 - t/(\sqrt{n} - \sqrt{m})] \leq e^{-t^2/2},$$

or equivalently, by setting $t = \epsilon(\sqrt{n} - \sqrt{m})$ for any $\epsilon \geq 0$,

$$\mathbb{P}[\sigma_{\min}(\mathbf{A}^\top)/(\sqrt{n} - \sqrt{m}) \leq 1 - \epsilon] \leq e^{-\epsilon^2(\sqrt{n} - \sqrt{m})^2/2}. \quad (40)$$

Let $\mathbf{p} \stackrel{\text{def}}{=} \mathbf{y} - \mathbf{A}\mathbf{x}^*$. The goal is to show that $\mathbf{p} = \mathbf{0}$. To this end, let $\mathbf{w} \stackrel{\text{def}}{=} \mathbf{A}^\dagger \mathbf{p}$, where $\mathbf{A}^\dagger = \mathbf{A}^\top(\mathbf{A}\mathbf{A}^\top)^{-1}$ denotes the MPP (note that $\mathbf{A}\mathbf{A}^\top$ is almost surely invertible). Because

$$\|\mathbf{w}\|^2 = \mathbf{p}^\top(\mathbf{A}\mathbf{A}^\top)^{-1}\mathbf{p} \leq \frac{\|\mathbf{p}\|^2}{\sigma_{\min}^2(\mathbf{A}^\top)}$$

we have for $0 < \epsilon < 1$, by the concentration of $\sigma_{\min}(\mathbf{A}^\top)$ (40) that with probability at least $1 - e^{-\epsilon^2(\sqrt{n} - \sqrt{m})^2/2}$,

$$\|\mathbf{w}\| \leq \frac{\|\mathbf{p}\|}{(\sqrt{n} - \sqrt{m})(1 - \epsilon)}.$$

Let $\mathbf{x}^\circ \stackrel{\text{def}}{=} \mathbf{x}^* + \mathbf{w}$ so that $\mathbf{y} - \mathbf{A}\mathbf{x}^\circ = \mathbf{0}$. Using this together with the optimality of \mathbf{x}^* gives

$$[\|\mathbf{y} - \mathbf{A}\mathbf{x}^*\| + \lambda f(\mathbf{x}^*)] - [\|\mathbf{y} - \mathbf{A}\mathbf{x}^\circ\| + \lambda f(\mathbf{x}^\circ)] \leq 0. \quad (41)$$

On the other hand,

$$\begin{aligned} [\|\mathbf{y} - \mathbf{A}\mathbf{x}^*\| + \lambda f(\mathbf{x}^*)] - [\|\mathbf{y} - \mathbf{A}\mathbf{x}^\circ\| + \lambda f(\mathbf{x}^\circ)] &= \|\mathbf{p}\| + \lambda f(\mathbf{x}^*) - \lambda f(\mathbf{x}^\circ) \\ &\geq \|\mathbf{p}\| - \lambda |f(\mathbf{x}^*) - \lambda f(\mathbf{x}^\circ)| \\ &\geq \|\mathbf{p}\| - \lambda L \|\mathbf{x}^* - \mathbf{x}^\circ\| \\ &= \|\mathbf{p}\| - \lambda L \|\mathbf{w}\| \\ &\geq \|\mathbf{p}\| \left(1 - \frac{\lambda L}{(\sqrt{n} - \sqrt{m})(1 - \epsilon)} \right), \end{aligned}$$

where, by (41), the last expression must be non-positive. But if we choose

⁷We actually use a one-sided variant of [Ver09, Corollary 5.35] which can be obtained by combining Lemma A.1(d) with the estimate of the expectation of σ_{\min} , [Ver09, Theorem 5.32].

$$\lambda < \frac{\sqrt{n} - \sqrt{m}}{L}(1 - \epsilon),$$

the only way to make it non-positive is that $\|\mathbf{p}\| = 0$. \square

Lemma B.2. Consider $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ and define

$$\begin{aligned} \Delta_p(\beta; \mathbf{h}, \lambda) &\stackrel{\text{def}}{=} \frac{1}{\sqrt{n}} \text{dist}(\beta \mathbf{h}, \|\cdot\|_{p^*} \leq \lambda) \\ \Delta_p(\beta; n, \lambda) &\stackrel{\text{def}}{=} \mathbb{E}[\Delta_p(\beta; \mathbf{h}, \lambda)] \end{aligned}$$

There exist universal constants $c_1, c_2 > 0$ such that for any $0 < \epsilon < 2$, any integers m, n , any $A > 0$ we have, with $\delta \stackrel{\text{def}}{=} (m-1)/n$ and $\phi(\alpha, \beta; \mathbf{g}, \mathbf{h})$ defined as in (30):

$$\mathbb{P}\{\exists 0 < \alpha \leq A, 0 < \beta < 1, |\phi(\alpha, \beta; \mathbf{g}, \mathbf{h}) - \kappa(\alpha, \beta)| \geq \epsilon\} \leq \frac{c_1}{\epsilon} e^{-c_2 \epsilon^2 n \cdot \frac{\min(\delta, 1/A^2)}{1+\delta A^2}} \stackrel{\text{def}}{=} \zeta(n, \epsilon; A, \delta) \quad (42)$$

where

$$\kappa(\alpha, \beta) = \kappa_p(\alpha, \beta; n, \delta, \lambda) \stackrel{\text{def}}{=} \beta \sqrt{\delta \alpha^2 + 1} - \alpha \Delta_p(\beta; n, \lambda). \quad (43)$$

Proof. We first look at the term $\|\alpha \mathbf{g}/\sqrt{n} - \mathbf{e}_1\|$. We have

$$\|\alpha \mathbf{g}/\sqrt{n} - \mathbf{e}_1\|^2 = \|\alpha \tilde{\mathbf{g}}/\sqrt{n}\|^2 + (\alpha g_0/\sqrt{n} - 1)^2,$$

where we partitioned \mathbf{g} as $[g_0, \tilde{\mathbf{g}}^\top]^\top$, $\tilde{\mathbf{g}} \in \mathbb{R}^{m-1}$. By Lemma A.1(b), for a standard Gaussian random vector $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$ it holds that for any $0 < \epsilon < 1$

$$\mathbb{P}\left\{\|x\|^2 \notin [(1-\epsilon)n, n/(1-\epsilon)]\right\} \leq 2e^{-\frac{\epsilon^2 n}{4}}. \quad (44)$$

Denoting $\delta \stackrel{\text{def}}{=} \frac{m-1}{n}$, it follows that

$$\mathbb{P}\left[\exists 0 < \alpha \leq A, \|\alpha \tilde{\mathbf{g}}/\sqrt{n}\|^2 \notin [(1-\epsilon)\delta \alpha^2, \delta \alpha^2/(1-\epsilon)]\right] \leq 2e^{-\frac{\epsilon^2(m-1)}{4}} = 2e^{-\frac{\epsilon^2 n \delta}{4}}.$$

Next, we show that the term $(\alpha g_0/\sqrt{n} - 1)^2$ cannot deviate much from 1: setting $\epsilon' = \epsilon/2$, we have $\sqrt{1-\epsilon} < 1-\epsilon'$ and $1+\epsilon' < 1/\sqrt{1-\epsilon}$ hence

$$\begin{aligned} \mathbb{P}\left\{\forall 0 < \alpha \leq A, \left(\frac{\alpha g_0}{\sqrt{n}} - 1\right)^2 \in [1-\epsilon, 1/(1-\epsilon)]\right\} &\geq \mathbb{P}\left\{\forall 0 < \alpha \leq A, \left|\frac{\alpha g_0}{\sqrt{n}} - 1\right| \in [1-\epsilon', 1+\epsilon']\right\} \\ &\geq \mathbb{P}\left\{\forall 0 < \alpha \leq A, \alpha g_0/\sqrt{n} \in [-\epsilon', \epsilon']\right\} \\ &= \mathbb{P}\{A g_0/\sqrt{n} \in [-\epsilon', \epsilon']\} \end{aligned}$$

so that using Lemma A.2 we have, with erfc the complementary error function,

$$\begin{aligned} \mathbb{P}\left\{\exists 0 < \alpha \leq A, \left(\frac{\alpha g_0}{\sqrt{n}} - 1\right)^2 \notin [1-\epsilon, 1/(1-\epsilon)]\right\} &\leq \mathbb{P}(|g_0| > \frac{\epsilon' \sqrt{n}}{A}) = \text{erfc}\left[\frac{\epsilon' \sqrt{n}}{\sqrt{2}A}\right] \\ &\leq e^{-\frac{\epsilon'^2 n}{2A^2}} = e^{-\frac{\epsilon^2 n}{8A^2}}. \end{aligned}$$

Combining the above, we get that: for any $0 < \epsilon < 1$, setting $\epsilon' = 1 - (1 - \epsilon)^2 = \epsilon(2 - \epsilon) \geq \epsilon$,

$$\begin{aligned} & \mathbb{P} \left\{ \exists 0 < \alpha \leq A, \quad \|\alpha \mathbf{g} / \sqrt{n} - \mathbf{e}_1\| \notin [(1 - \epsilon) \sqrt{\delta \alpha^2 + 1}, \sqrt{\delta \alpha^2 + 1} / (1 - \epsilon)] \right\} \\ &= \mathbb{P} \left\{ \exists 0 < \alpha \leq A, \quad \|\alpha \mathbf{g} / \sqrt{n} - \mathbf{e}_1\|^2 \notin [(1 - \epsilon')(\delta \alpha^2 + 1), (\delta \alpha^2 + 1) / (1 - \epsilon')] \right\} \\ &\leq 2e^{-\frac{\epsilon'^2 n \delta}{4}} + e^{-\frac{\epsilon'^2 n}{8A^2}} \leq c'_1 e^{-c'_2 \epsilon'^2 n \cdot \min(\delta, 1/A^2)} \leq c'_1 e^{-c'_2 \epsilon^2 n \cdot \min(\delta, 1/A^2)}, \end{aligned} \quad (45)$$

where the constants c'_1, c'_2 are universal.

For the second term in $\phi(\alpha, \beta; \mathbf{g}, \mathbf{h})$, we use Gaussian concentration of Lipschitz functions. It is known that (Euclidean) distance to a convex set, $\text{dist}(\mathbf{x}, \mathcal{C})$, is 1-Lipschitz in \mathbf{x} (with respect to the Euclidean metric) so by Lemma A.1(e) we get for $0 < \beta \leq 1$ (we omit the dependency in λ for brevity) and any $t \geq 0$

$$\mathbb{P} \left[\frac{\sqrt{n}}{\beta} |\Delta_p(\beta; \mathbf{h}) - \Delta_p(\beta)| > t \right] \leq 2e^{-\frac{t^2}{2}},$$

Setting $t = \tau \sqrt{n} / \beta$ we obtain for any $\tau \geq 0$

$$\mathbb{P} [|\Delta_p(\beta; \mathbf{h}) - \Delta_p(\beta)| > \tau] \leq 2e^{-\frac{\tau^2 n}{2\beta^2}} \leq 2e^{-\frac{\tau^2 n}{2}}, \quad (46)$$

This obviously extends to $\beta = 0$ since $\Delta_p(0; \mathbf{h}) = \Delta_p(0) = 0$.

Next we want to bound

$$\mathbb{P} [\exists \beta \in [0, 1], |\Delta_p(\beta; \mathbf{h}) - \Delta_p(\beta)| > \tau] = \mathbb{P} \left[\sup_{\beta \in [0, 1]} |\Delta_p(\beta; \mathbf{h}) - \Delta_p(\beta)| > \tau \right].$$

By Lemma B.13, the function $\beta \mapsto f_{\mathbf{h}}(\beta) \stackrel{\text{def}}{=} |\Delta_p(\beta; \mathbf{h}) - \Delta_p(\beta)|$ is $L_{\mathbf{h}}$ -Lipschitz in β with $L_{\mathbf{h}} \stackrel{\text{def}}{=} \max \{\|\mathbf{h}\| / \sqrt{n}, 1\}$. Hence $f_{\mathbf{h}}$ is continuous, and its supremum on the closed interval $[0, 1]$ is indeed a maximum reached at some maximizer $\beta_{\mathbf{h}}$.

Let $b \leq 1/2$ and $Y_b = \{b\tau/2, 3b\tau/2, \dots, (n-1/2)b\tau\}$ be a uniform sampling of $[0, 1]$ with spacing $b\tau$, with the last segment possibly being shorter. For a given \mathbf{h} , there exists $y_{\mathbf{h}} \in Y_b$ such that $|\beta_{\mathbf{h}} - y_{\mathbf{h}}| \leq b\tau$. For this $y_{\mathbf{h}}$ we write

$$f_{\mathbf{h}}(\beta_{\mathbf{h}}) - f_{\mathbf{h}}(y_{\mathbf{h}}) \leq |f_{\mathbf{h}}(\beta_{\mathbf{h}}) - f_{\mathbf{h}}(y_{\mathbf{h}})| \leq L_{\mathbf{h}} |\beta_{\mathbf{h}} - y_{\mathbf{h}}| \leq L_{\mathbf{h}} b\tau$$

so that

$$\begin{aligned} \mathbb{P} \left[\sup_{\beta \in [0, 1]} f_{\mathbf{h}}(\beta) > \tau \right] &= \mathbb{P} [f_{\mathbf{h}}(\beta_{\mathbf{h}}) > \tau] \leq \mathbb{P} [f_{\mathbf{h}}(y_{\mathbf{h}}) + L_{\mathbf{h}} b\tau > \tau] \leq \mathbb{P} \left[f_{\mathbf{h}}(y_{\mathbf{h}}) > \tau/2 \vee L_{\mathbf{h}} b\tau > \tau/2 \right] \\ &\leq \mathbb{P} [f_{\mathbf{h}}(y_{\mathbf{h}}) > \tau/2] + \mathbb{P} [\max \{\|\mathbf{h}\| / \sqrt{n}, 1\} b > 1/2] \end{aligned}$$

As we do not know a priori to which $y \in Y_b$ the maximizer $\beta_{\mathbf{h}}$ will be close, we continue with a union bound, and we further use that $b \leq 1/2$ to obtain by (46) and Lemma A.1(a)

$$\begin{aligned} \mathbb{P} \left[\sup_{\beta \in [0, 1]} f_{\mathbf{h}}(\beta) > \tau \right] &\leq \mathbb{P} \left[\bigvee_{y \in Y_b} f_{\mathbf{h}}(y) > \tau/2 \right] + \mathbb{P} \left[\|\mathbf{h}\|^2 > \frac{n}{4b^2} \right] \leq \#Y_b 2e^{-\frac{\tau^2 n}{8}} + e^{-\frac{(1-4b^2)^2 n}{4}} \\ &\leq (1 + 1/(b\tau)) 2e^{-\frac{\tau^2 n}{8}} + e^{-\frac{(1-4b^2)^2 n}{4}}. \end{aligned}$$

We conclude the argument for the second term by choosing $b \stackrel{\text{def}}{=} \frac{1}{2} \sqrt{\frac{2-\sqrt{2}}{2}} \approx 0.27$, so that: for any $0 < \tau < 1$, we have $\frac{(1-4b^2)^2}{4} = \frac{1}{8} \geq \frac{\tau^2}{8}$ hence

$$\mathbb{P}[\sup_{\beta \in [0,1]} f_{\mathbf{h}}(\beta) > \tau] \leq (3 + 2/(b\tau))e^{-\frac{\tau^2 n}{8}} < \frac{11}{\tau} e^{-\frac{\tau^2 n}{8}}, \quad 0 < \tau < 1. \quad (47)$$

To conclude we combine concentration bounds for both terms. First, we observe that

$$\sup_{\substack{0 \leq \alpha \leq A \\ 0 \leq \beta \leq 1}} |\phi(\alpha, \beta; \mathbf{g}, \mathbf{h}) - \kappa(\alpha, \beta)| \leq \sup_{0 \leq \alpha \leq A} \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \sqrt{1 + \delta \alpha^2} + A \sup_{0 \leq \beta \leq 1} |\Delta_p(\beta; \mathbf{h}) - \Delta_p(\beta)|$$

hence by a union bound we just need to control the probability that each term exceeds $\epsilon/2$. Since we assume that $0 < \epsilon < 2$, we can use the multiplicative control (45) of the first term and convert it into the desired additive bound as follows:

$$\begin{aligned} \mathbb{P}\left\{ \exists \alpha \in (0, A], \left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\| - \sqrt{\delta \alpha^2 + 1} > \frac{\epsilon}{2} \right\} &= \mathbb{P}\left\{ \exists \alpha \in (0, A], \left| \frac{\left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\|}{\sqrt{\delta \alpha^2 + 1}} - 1 \right| > \frac{\epsilon}{2\sqrt{\delta \alpha^2 + 1}} \right\} \\ &\leq \mathbb{P}\left\{ \exists \alpha \in (0, A], \left| \frac{\left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\|}{\sqrt{\delta \alpha^2 + 1}} - 1 \right| > \frac{\epsilon}{2\sqrt{\delta A^2 + 1}} \right\} \\ &\leq \mathbb{P}\left\{ \exists \alpha \in (0, A], \frac{\left\| \frac{\alpha \mathbf{g}}{\sqrt{n}} - \mathbf{e}_1 \right\|}{\sqrt{\delta \alpha^2 + 1}} \notin [1 - \epsilon', \frac{1}{1 - \epsilon'}] \right\} \\ &\leq c'_1 e^{-c'_2 \epsilon'^2 n \cdot \min(\delta, 1/A^2)}. \end{aligned}$$

provided that $[1 - \epsilon', 1/(1 - \epsilon')] \subset [1 - \frac{\epsilon}{2\sqrt{\delta A^2 + 1}}, 1 + \frac{\epsilon}{2\sqrt{\delta A^2 + 1}}]$. This is achieved with $\epsilon' = \epsilon/(1 + \sqrt{\delta A^2 + 1})$. Combining the resulting probability bound with the bound on $\mathbb{P}[\sup_{\beta \in [0,1]} f_{\mathbf{h}}(\beta) > \frac{\epsilon}{2A}]$ resulting from (47) yields the result. \square

Corollary B.1. Consider $A > 0$ and define for any set $S \subseteq [0, A]$:

$$\begin{aligned} \phi(\alpha; \mathbf{g}, \mathbf{h}) &\stackrel{\text{def}}{=} \sup_{0 \leq \beta \leq 1} \phi(\alpha, \beta; \mathbf{g}, \mathbf{h}), & \kappa(\alpha) &\stackrel{\text{def}}{=} \sup_{0 \leq \beta \leq 1} \kappa(\alpha, \beta), \\ \phi_S(\mathbf{g}, \mathbf{h}) &\stackrel{\text{def}}{=} \inf_{\alpha \in S} \phi(\alpha; \mathbf{g}, \mathbf{h}), & \kappa_S &\stackrel{\text{def}}{=} \inf_{\alpha \in S} \kappa(\alpha), \end{aligned}$$

With $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $\delta = (m - 1)/n$, we have for $0 < \epsilon < 2$

$$\mathbb{P}\left[\sup_{S \subset [0, A]} |\phi_S(\mathbf{g}, \mathbf{h}) - \kappa_S| \geq \epsilon \right] \leq \zeta(n, \epsilon). \quad (48)$$

with $\zeta(n, \epsilon) = \zeta(n, \epsilon; A, \delta)$ defined in (42). In particular,

$$\mathbb{P}\left[\sup_{0 \leq \alpha \leq A} |\phi(\alpha) - \kappa(\alpha)| \geq \epsilon \right] \leq \zeta(n, \epsilon). \quad (49)$$

Proof. To lighten notation we suppress the dependence of the stochastic function ϕ on random vectors \mathbf{g} and \mathbf{h} . By Lemma B.2 we have with probability at least $1 - \zeta(n, \epsilon)$: for all $0 \leq \alpha \leq A$

and $0 \leq \beta \leq 1$, $|\phi(\alpha, \beta) - \kappa(\alpha, \beta)| \leq \epsilon$. When this holds we have for any $S \subset [0, A]$:

$$\phi_S = \inf_{\alpha \in S} \phi(\alpha) \leq \inf_{\alpha \in S} [\kappa(\alpha) + \epsilon] = \kappa_S + \epsilon \quad (50)$$

and

$$\phi_S = \inf_{\alpha \in S} \phi(\alpha) \geq \inf_{\alpha \in S} [\kappa(\alpha) - \epsilon] = \kappa_S - \epsilon. \quad (51)$$

□

We will shortly characterize $\kappa(\alpha, \beta)$ and $\kappa(\alpha)$. To that end, we will use some properties of the following quantity

$$D_p(t; n) \stackrel{\text{def}}{=} \left(\mathbb{E} \left[\frac{1}{\sqrt{n}} \text{dist}(\mathbf{h}, \|\cdot\|_{p^*} \leq t) \right] \right)^2. \quad (52)$$

Lemma B.3 (Deterministic properties of D_p). Define $\mathcal{C}_t = \mathcal{C}_{t,p} \stackrel{\text{def}}{=} \{ \mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_{p^*} \leq t \}$ and $D_p(t; n)$ as in (52). Using $D_p(t)$ as a shorthand, the following hold:

1. The sets \mathcal{C}_t are convex and nested with $\mathcal{C}_t \subsetneq \mathcal{C}_{t'}$ for $t < t'$.
2. For any vector \mathbf{h} , the function $t \mapsto \text{dist}(\mathbf{h}, \mathcal{C}_t)$ is non-increasing and convex.
3. $D_p(t)$ is a (strictly) decreasing convex function of t ,
4. $\lim_{t \rightarrow \infty} D_p(t) = 0$,
5. $\frac{n}{n+1} \leq D_p(0) \leq 1$,
6. The function $t \mapsto D_p(t)$ is infinitely differentiable.
7. Let $g(t) = g(t; n) \stackrel{\text{def}}{=} D_p(t) - \frac{t}{2} D'_p(t)$. For any $0 < \delta < D_p(0)$ there is a unique

$$t^* = t_p^*(\delta; n) \in (0, \infty)$$

such that $g(t) > \delta$ for $t < t^*$ and $g(t) < \delta$ for $t > t^*$. It holds that $D_p(t^*) < g(t^*) = \delta$.

Proof. • Property 1 is obvious.

- Property 2: we recall that $\mathbf{x} \mapsto \text{dist}(\mathbf{x}, \mathcal{C})$ is convex for convex \mathcal{C} [BV04, Example 3.16]. Next, $(\mathbf{x}, t) \mapsto t \text{dist}(\mathbf{x}/t, \mathcal{C})$ is convex in both arguments because it is the perspective of $\mathbf{x} \mapsto \text{dist}(\mathbf{x}, \mathcal{C})$ [BV04, Chapter 2]. Applying this to $\mathcal{C} = \mathcal{C}_1$ and observing that $\text{dist}(\mathbf{h}, \mathcal{C}_t) = t \text{dist}(\mathbf{h}/t, \mathcal{C}_1)$ we obtain that $\text{dist}(\mathbf{h}, \mathcal{C}_t)$ is convex in t . The fact that it is non-increasing follows from Property 1.
- Property 3: since expectation of convex functions is convex, and the pointwise square of a non-negative convex function is convex, we conclude that $D_p(t)$ is convex as claimed. The fact that it is (strictly) decreasing is obvious.
- Property 4: for any $\mathbf{y} \in \mathbb{R}^n$ and $p \geq 1$, $\|\mathbf{y}\|_{p^*} \leq \|\mathbf{y}\|_1 \leq \sqrt{n} \|\mathbf{y}\|$. Hence, for any given

$t > 0$, $\text{dist}(\mathbf{y}, \mathcal{C}_t) = 0$ as soon as $\|\mathbf{y}\| \leq t/\sqrt{n}$. Further, for all $p \geq 1$ we can write

$$\begin{aligned}
\mathbb{E}[\text{dist}(\mathbf{h}, \mathcal{C}_t)] &= \int_{\mathbb{R}^n} \text{dist}(\mathbf{y}, \mathcal{C}_t) p_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} \\
&= \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{r=0}^{\infty} \text{dist}(r\mathbf{b}, \mathcal{C}_t) p_{\mathbf{h}}(r\mathbf{b}) \mu(d\mathbf{b}) r^{n-1} dr \\
&= \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{r=t/\sqrt{n}}^{\infty} \text{dist}(r\mathbf{b}, \mathcal{C}_t) p_{\mathbf{h}}(r\mathbf{b}) \mu(d\mathbf{b}) r^{n-1} dr \\
&\leq \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{r=t/\sqrt{n}}^{\infty} r p_{\mathbf{h}}(r\mathbf{b}) \mu(d\mathbf{b}) r^{n-1} dr \\
&= \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \mu(d\mathbf{b}) \int_{r=t/\sqrt{n}}^{\infty} r^n Z_n e^{-r^2/2} dr \\
&= Z_n \mu(\mathbb{S}^{n-1}) \int_{r=t/\sqrt{n}}^{\infty} r^n e^{-r^2/2} dr \rightarrow 0 \text{ as } t \rightarrow \infty,
\end{aligned}$$

where $Z_n = (2\pi)^{-n/2}$ is the normalization term for the n -variate iid Gaussian distribution, and $\mu(\mathbb{S}^{n-1}) = 2\pi^{n/2}/\Gamma(\frac{n}{2})$ is the surface area of the unit-radius $(n-1)$ -sphere embedded in \mathbb{R}^n with $\Gamma(\cdot)$ being the gamma function.

- Property 5: by Jensen's inequality and Property 3, we obtain the upper bound

$$D_p(0) = \left(\mathbb{E} \frac{1}{\sqrt{n}} \text{dist}(\mathbf{h}, \mathcal{C}_0) \right)^2 \leq \mathbb{E} \left(\frac{1}{\sqrt{n}} \text{dist}(\mathbf{h}, \mathcal{C}_0) \right)^2 = \mathbb{E} \frac{\|\mathbf{h}\|^2}{n} = 1.$$

To get the lower bound, we again note that $nD_p(0) = [\mathbb{E} \|\mathbf{h}\|]^2$, which can be computed by integration in polar coordinates. Working as in the proof of Property 4 we get that

$$\mathbb{E}[\text{dist}(\mathbf{h}, \mathcal{C}_0)] = \mathbb{E}[\|\mathbf{h}\|] = \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{r=0}^{\infty} r p_{\mathbf{h}}(r\mathbf{b}) \mu(d\mathbf{b}) r^{n-1} dr$$

since $r = \|\mathbf{h}\| = \text{dist}(\mathbf{h}, \mathcal{C}_0)$. It follows that

$$\begin{aligned}
\mathbb{E}[\|\mathbf{h}\|] &= Z_n \mu(\mathbb{S}^{n-1}) \int_{r=0}^{\infty} r^n e^{-r^2/2} dr \\
&\stackrel{(a)}{=} (2\pi)^{-n/2} \frac{2\pi^{n/2}}{\Gamma(\frac{n}{2})} \int_{s=0}^{\infty} 2^{\frac{n-1}{2}} s^{\frac{n-1}{2}} e^{-s} ds \\
&\stackrel{(b)}{=} \sqrt{2} \Gamma\left(\frac{n+1}{2}\right) / \Gamma\left(\frac{n}{2}\right)
\end{aligned}$$

where in (a) we used the substitution $u = r^2/2$, and in (b) we invoked the definition of the gamma function, $\Gamma(z) = \int_0^{\infty} x^{z-1} e^{-x} dx$. We now use the inequality of Wendel, [Wen48, Eq. (7)]:

$$\Gamma(x+a)/\Gamma(x) \geq x(x+a)^{a-1}$$

to conclude that

$$\mathbb{E}[\|\mathbf{h}\|] \geq n(n+1)^{-1/2} \tag{53}$$

and $D_p(0) \geq n/(n+1)$.

- Property 6: starting similarly as in the proof of Property 4 and using $\text{dist}(\mathbf{h}, \mathcal{C}_t) = t \text{dist}(\mathbf{h}/t, \mathcal{C}_1)$ and a change of variable $r = t\rho$ we obtain the following expression for $D_p(t)$:

$$q(t) \stackrel{\text{def}}{=} \sqrt{nD_p(t)} = t^{n+1} \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{\rho=0}^{\infty} \text{dist}(\rho\mathbf{b}, \mathcal{C}_1) p_{\mathbf{h}}(\rho t\mathbf{b}) \mu(d\mathbf{b}) \rho^{n-1} d\rho \quad (54)$$

$$= Z_n t^{n+1} \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{\rho=0}^{\infty} \text{dist}(\rho\mathbf{b}, \mathcal{C}_1) e^{-\rho^2 t^2/2} \mu(d\mathbf{b}) \rho^{n-1} d\rho \quad (55)$$

so that $D_p(t)$ is infinitely differentiable (by the repeated application of the dominated convergence theorem).

- Property 7: in particular, since

$$\left| \frac{\partial}{\partial t} \text{dist}(\rho\mathbf{b}, \mathcal{C}_1) e^{-\rho^2 t^2/2} \rho^{n-1} \right| = \text{dist}(\rho\mathbf{b}, \mathcal{C}_1) t \rho^{n+1} e^{-\rho^2 t^2/2} \leq \rho^{n+2} t e^{-\rho^2 t^2/2}$$

where the rightmost expression is integrable for every $t > 0$ and $n \in \mathbb{N}$, the dominated convergence theorem allows us to differentiate under the integral sign in (55) to get

$$\begin{aligned} -\frac{t}{2} D'_p(t) &= -\frac{t}{n} q(t) q'(t) \\ &= -\frac{1}{n} q(t) \left((n+1)q(t) - Z_n t^{n+3} \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{\rho=0}^{\infty} \text{dist}(\rho\mathbf{b}, \mathcal{C}_1) \rho^{n+1} e^{-\rho^2 t^2/2} \mu(d\mathbf{b}) d\rho \right). \end{aligned} \quad (56)$$

All terms can be seen to vanish as $t \rightarrow \infty$ by arguments analogous to those in the end of the proof of Property 5, hence $\lim_{t \rightarrow \infty} [-\frac{t}{2} D'_p(t)] = 0$.

Since $D_p(t)$ is strictly decreasing we have $D'_p(t) < 0$. Since it is convex, we have $D''_p(t) \geq 0$. Thus

$$g'(t) = D'_p(t) - \frac{1}{2} D'_p(t) - \frac{t}{2} D''_p(t) = \frac{1}{2} D'_p(t) - \frac{t}{2} D''_p(t) < 0$$

for $t > 0$, meaning that $g(t)$ is strictly decreasing. Since $\lim_{t \rightarrow \infty} [-\frac{t}{2} D'_p(t)] = 0$ and $\lim_{t \rightarrow \infty} D_p(t) = 0$, we have $\lim_{t \rightarrow \infty} g(t) = 0$, it follows that for $0 < t < \lim_{t \rightarrow 0} g(t)$, there is a unique $t^*(\delta)$, such that $0 < t^*(\delta) < \infty$ and $g(t) > \delta$ for $t < t^*$ and $g(t) < \delta$ for $t > t^*$. We conclude by observing that since $g(t) \geq D_p(t)$ we have $\lim_{t \rightarrow 0} g(t) \geq D_p(0)$. \square

Lemma B.4. Denote

$$\theta(t) \stackrel{\text{def}}{=} \mathbb{E}(|h| - t)_+^2 \quad (57)$$

with h a standard centered normal variable and $(\cdot)_+ = \max(\cdot, 0)$ the positive part. This is a strictly decreasing function of t with $\theta(0) = 1$ and $\lim_{t \rightarrow \infty} \theta(t) = 0$. The following holds for all $1 \leq p \leq \infty$, $0 < \delta < 1$:

1. For any $n \geq 1$

$$t_p^*(\delta; n) \leq 2 \theta^{-1}(\delta) n^{1-1/p}. \quad (58)$$

2. For any $n \geq \frac{2}{1-\delta}$

$$t_p^*(\delta; n) \geq \theta^{-1}\left(\frac{1+\delta}{2}\right) > 0 \quad (59)$$

3. For any $n \geq \max\left(\frac{2}{1-\delta}, \frac{1}{\delta}\right)$

$$t_p^*(\delta; n) \geq \max\left(\theta^{-1}(2\delta), \theta^{-1}\left(\frac{1+\delta}{2}\right)\right) > 0 \quad (60)$$

4. There is a universal constant C independent of δ , p and n such that for all $n \geq 1$

$$D_p(t_p^*; n) \geq \left(\frac{\delta}{C}\right)^2 \quad (61)$$

where we use the shorthand $t_p^* = t_p^*(\delta; n)$.

Lemma B.5. With the notations of Theorem 3.1, for $p = 2$,

$$\lim_{n \rightarrow \infty} -t_2^* D_2'(t_2^*; n) = (1 - \delta)\delta > 0 \quad (62)$$

$$\lim_{n \rightarrow \infty} \alpha_2^*(\delta; n) = \frac{1}{\sqrt{1-\delta}}. \quad (63)$$

Lemma B.6. With the notations of Theorem 3.1, for $p = 1$,

$$\lim_{n \rightarrow \infty} -t_1^* D_1'(t_1^*; n) > 0 \quad (64)$$

$$\alpha_1^*(\delta) \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} \alpha_1^*(\delta; n) = \sqrt{\frac{1}{\sqrt{\frac{2}{\pi}} e^{-\frac{(t_1^*)^2}{2}} t_1^* - \delta (t_1^*)^2} - \frac{1}{\delta}}. \quad (65)$$

where $t_1^*(\delta) \stackrel{\text{def}}{=} \sqrt{2} \operatorname{erfc}^{-1}(\delta)$.

Proof of Lemmas B.4-B.5-B.6.

- **Step 1.** We prove that for all p, n, t

$$D_p(t; n) \geq D_1(t; n) \geq D_p\left(tn^{1-1/p}; n\right). \quad (66)$$

The inequalities $\|\cdot\|_{p^*} \geq \|\cdot\|_\infty \geq n^{-1/p^*} \|\cdot\|_{p^*} = n^{-(1-1/p)} \|\cdot\|_{p^*}$ imply the inclusions $\mathcal{C}_{t,p} \subset \mathcal{C}_{t,1} \subset \mathcal{C}_{tn^{1-1/p}, p}$. It follows that $d(\cdot, \mathcal{C}_{t,p}) \geq d(\cdot, \mathcal{C}_{t,1}) \geq d(\cdot, \mathcal{C}_{tn^{1-1/p}, p})$, which yields (66).

- **Step 2.** We establish that for any p, n, δ , with the shorthand $t_p^* = t_p^*(\delta; n)$,

$$D_p(t_p^*/2; n) \geq \delta \quad (67)$$

$$D_p(t_p^*; n) \leq \delta \quad (68)$$

With the additional shorthand $D_p(t) = D_p(t; n)$, since $D_p(t)$ is convex, we have for all t, h

$$D_p(t+h) \geq D_p(t) + hD_p'(t). \quad (69)$$

Applying it to $t = t_p^*$ and $h = -t_p^*/2$ and using the definition of t_p^* , we get

$$D_p(t_p^*/2) \geq D_p(t_p^*) - (t_p^*/2)D_p'(t_p^*) = \delta$$

i.e. (67) holds. Since $D_p(t)$ is non-increasing, we have $D_p'(t) \leq 0$ and $D_p(t) - \frac{t}{2}D_p'(t) \geq D_p(t)$ for any t . Applying to $t = t_p^*$ this establishes (68) by definition of t_p^* .

- **Step 3.** We show that for any p, n, t

$$D_p(t; n) \leq \bar{D}_p(t; n) \stackrel{\text{def}}{=} \frac{1}{n} \mathbb{E} \text{dist}^2(\mathbf{h}, \|\cdot\|_{p^*} \leq t) \quad (70)$$

and that for $p = 1$

$$\bar{D}_1(t; n) = \theta(t) \quad (71)$$

is independent of n . By the concavity of square root and Jensen's inequality we can write

$$nD_p(t; n) = \left(\mathbb{E} \text{dist}(\mathbf{h}, \|\cdot\|_{p^*} \leq t) \right)^2 \leq \mathbb{E} \text{dist}^2(\mathbf{h}, \|\cdot\|_{p^*} \leq t) = n\bar{D}_p(t; n).$$

This establishes (70).

For $p = 1$, since $p^* = \infty$ we can compute

$$\begin{aligned} n\bar{D}_1(t) &\stackrel{\text{def}}{=} \mathbb{E} \text{dist}^2(\mathbf{h}, \|\cdot\|_{\infty} \leq t) = \mathbb{E} \left\| \mathbf{h} - \text{proj}_{\|\cdot\|_{\infty} \leq t} \mathbf{h} \right\|_2^2 = \mathbb{E} \sum_{i=1}^n (h_i - (\text{proj}_{\|\cdot\|_{\infty} \leq t} \mathbf{h})_i)^2 \\ &= \sum_{i=1}^n \mathbb{E} (|h_i| - t)_+^2 = n\theta(t). \end{aligned}$$

- **Step 4.** Combining the previous steps we get for any p, δ, n :

$$\theta(t_p^* n^{1/p-1}/2) \stackrel{(71)}{=} \bar{D}_1(t_p^* n^{1/p-1}/2) \stackrel{(70)}{\geq} D_1(t_p^* n^{1/p-1}/2) \stackrel{(66)}{\geq} D_p(t_p^*/2) \stackrel{(67)}{\geq} \delta$$

This yields (58).

- **Step 5.** We show that for any p, n, t

$$D_p(t; n) \geq \bar{D}_p(t; n) - \frac{1}{n}. \quad (72)$$

Indeed, the function $f : \mathbf{h} \mapsto f(\mathbf{h}) \stackrel{\text{def}}{=} \text{dist}(\mathbf{h}, \|\cdot\|_{p^*} \leq t)$ is 1-Lipschitz, and \mathbf{h} is a standard normal Gaussian variable, hence we can apply Lemma A.1(f).

- **Step 6.** Combining with the previous steps yields for any p, δ , and $n > \frac{1}{1-\delta}$:

$$\theta(t_p^*) \stackrel{(71)}{=} \bar{D}_1(t_p^*) \stackrel{(72)}{\leq} D_1(t_p^*) + \frac{1}{n} \stackrel{(66)}{\leq} D_p(t_p^*) + \frac{1}{n} \stackrel{(68)}{\leq} \delta + \frac{1}{n} < 1$$

For $n \geq \frac{2}{1-\delta}$, we have $\delta + \frac{1}{n} \leq \frac{1+\delta}{2} < 1$ which yields (59). For $n \geq \max(\frac{2}{1-\delta}, 1/\delta)$, we have $\delta + \frac{1}{n} \leq \min(\frac{1+\delta}{2}, 2\delta) < 1$ which yields (60).

- **Step 7.** To establish (61) we start with the expression (56), with the shorthand $q(t) \stackrel{\text{def}}{=} \sqrt{nD_p(t)}$:

$$\begin{aligned} -\frac{t}{2} D_p'(t) &= -\frac{1}{n} q(t) \left((n+1)q(t) - Z_n t^{n+3} \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{\rho=0}^{\infty} \text{dist}(\rho \mathbf{b}, \mathcal{C}_1) \rho^{n+1} e^{-\rho^2 t^2/2} \mu(d\mathbf{b}) d\rho \right) \\ -\frac{t}{2} D_p'(t) + \frac{n+1}{n} q^2(t) &= \frac{q(t)}{n} Z_n t^{n+3} \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{\rho=0}^{\infty} \text{dist}(\rho \mathbf{b}, \mathcal{C}_1) \rho^{n+1} e^{-\rho^2 t^2/2} \mu(d\mathbf{b}) d\rho \\ &\stackrel{r=\rho \mathbf{b}}{=} \frac{q(t)}{n} \int_{\mathbf{b} \in \mathbb{S}^{n-1}} \int_{\rho=0}^{\infty} t \text{dist}(r \mathbf{b}/t, \mathcal{C}_1) r^{n+1} p_{\mathbf{h}}(r \mathbf{b}) \mu(d\mathbf{b}) dr \\ &\stackrel{\mathbf{y}=r \mathbf{b}}{=} \frac{q(t)}{n} \int_{\mathbb{R}^n} t \text{dist}(\mathbf{y}/t, \mathcal{C}_1) \|\mathbf{y}\|^2 p_{\mathbf{h}}(\mathbf{y}) d\mathbf{y} \\ &= \frac{q(t)}{n} \mathbb{E}[\text{dist}(\mathbf{h}, \mathcal{C}_t) \|\mathbf{h}\|^2]. \end{aligned} \quad (73)$$

Since $\frac{n+1}{n}q^2(t) = (n+1)D_p(t)$ we can rewrite (73) using the definition of $D_p(t)$ as

$$D_p(t) - \frac{t}{2}D'_p(t) = \frac{q(t)}{n}\mathbb{E}[\text{dist}(\mathbf{h}, \mathcal{C}_t)(\|\mathbf{h}\|^2 - n)]. \quad (74)$$

Observing further that $\mathbb{E}[\|\mathbf{h}\|^2 - n] = 0$ and $\mathbb{E}[\text{dist}(\mathbf{h}, \mathcal{C}_t)] = q(t)$, the following holds:

$$\begin{aligned} D_p(t) - \frac{t}{2}D'_p(t) &= \frac{q(t)}{n} \mathbb{E} \left[(\text{dist}(\mathbf{h}, \mathcal{C}_t) - q(t))(\|\mathbf{h}\|^2 - n) \right] \\ &= \sqrt{D_p(t)} \mathbb{E} \left[(\text{dist}(\mathbf{h}, \mathcal{C}_t) - q(t)) \frac{\|\mathbf{h}\|^2 - n}{\sqrt{n}} \right] \\ &\leq \sqrt{D_p(t)} \mathbb{E} \left[\left| \text{dist}(\mathbf{h}, \mathcal{C}_t) - q(t) \right| \left| \frac{\|\mathbf{h}\|^2 - n}{\sqrt{n}} \right| \right] \\ &= \sqrt{D_p(t)} \int_0^\infty \mathbb{P} \left[\left| \text{dist}(\mathbf{h}, \mathcal{C}_t) - q(t) \right| \left| \frac{\|\mathbf{h}\|^2 - n}{\sqrt{n}} \right| \geq \epsilon \right] d\epsilon \end{aligned}$$

The integrand can be controlled by a union bound as

$$\mathbb{P} \left[\left| \text{dist}(\mathbf{h}, \mathcal{C}_t) - q(t) \right| \left| \frac{\|\mathbf{h}\|^2 - n}{\sqrt{n}} \right| \geq \epsilon \right] \leq \mathbb{P} \left[\left| \text{dist}(\mathbf{h}, \mathcal{C}_t) - q(t) \right| \geq \sqrt{\epsilon} \right] + \mathbb{P} \left[\left| \frac{\|\mathbf{h}\|^2 - n}{\sqrt{n}} \right| \geq \sqrt{\epsilon} \right]$$

which together with Lemma A.1(e) for the first term and Lemma A.1(c) for the second term yields $D_p(t) - \frac{t}{2}D'_p(t) \leq C_n \sqrt{D_p(t)}$ with

$$\begin{aligned} C_n &\stackrel{\text{def}}{=} \int_0^\infty 2 \left(e^{-\epsilon/2} + \begin{cases} e^{-\frac{\epsilon}{8}} & \text{for } 0 \leq \epsilon \leq n, \\ e^{-\frac{\sqrt{\epsilon n}}{8}} & \text{for } \epsilon > n \end{cases} \right) d\epsilon \\ &\leq \int_0^\infty 2 \left(e^{-\epsilon/2} + e^{-\frac{\epsilon}{8}} + e^{-\frac{\sqrt{\epsilon n}}{8}} \right) d\epsilon \\ &\leq \int_0^\infty 2 \left(e^{-\epsilon/2} + e^{-\frac{\epsilon}{8}} + e^{-\frac{\sqrt{\epsilon}}{8}} \right) d\epsilon \stackrel{\text{def}}{=} C < \infty. \end{aligned}$$

- **Step 8. (Proof of Lemma B.5).** Since $\text{dist}(\mathbf{h}, \|\cdot\|_{p^*} \leq t) = (\|\mathbf{h}\|_2 - t)_+$, we have $D_2(t; n) = \frac{1}{n} (\mathbb{E}(\|\mathbf{h}\|_2 - t)_+)^2$. With $d(t) \stackrel{\text{def}}{=} \sqrt{D_2(t; n)} = \frac{1}{\sqrt{n}} \mathbb{E}(\|\mathbf{h}\|_2 - t)_+$ we have

$$d'(t) = -\frac{1}{\sqrt{n}} \mathbb{E}[\mathbb{I}_{\|\mathbf{h}\|_2 > t}] = -\frac{1}{\sqrt{n}} \mathbb{P}(\|\mathbf{h}\|_2 > t)$$

and $-(t/2)D'_2(t) = -td(t)d'(t)$. With a change of variables $\tau = t/\sqrt{n}$, define

$$\begin{aligned} F_1(\tau; n) &\stackrel{\text{def}}{=} D_2(\tau\sqrt{n}; n) = \left\{ \mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ \right\}^2 \\ F_2(\tau; n) &\stackrel{\text{def}}{=} -\frac{\tau\sqrt{n}}{2}D'_2(\tau\sqrt{n}; n) = \mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ \cdot \tau \mathbb{P} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} > \tau \right) \\ F(\tau; n) &\stackrel{\text{def}}{=} F_1(\tau; n) + F_2(\tau; n) \end{aligned}$$

Since $\|\mathbf{h}\|_2$ concentrates around \sqrt{n} for large n , as we show below for any $0 < \tau < 1$

$$\lim_{n \rightarrow \infty} F_1(\tau; n) \stackrel{\text{def}}{=} F_1(\tau) = (1 - \tau)^2 \quad (75)$$

$$\lim_{n \rightarrow \infty} F_2(\tau; n) \stackrel{\text{def}}{=} F_2(\tau) = (1 - \tau)\tau \quad (76)$$

It follows that for $\delta \in (0, 1)$ we have $\lim_{n \rightarrow \infty} F(1 - \delta; n) = \delta$ and

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{t_2^*(\delta; n)}{\sqrt{n}} &= 1 - \delta \\ \lim_{n \rightarrow \infty} -(t_2^*/2)D'_2(t_2^*; n) &= F_2(\delta) = (1 - \delta)\delta \\ \lim_{n \rightarrow \infty} D(t_2^*; n) &= F_1(\delta) = \delta^2 \\ \lim_{n \rightarrow \infty} \alpha_2^*(\delta; n) &= \sqrt{\frac{\delta^2}{\delta(\delta - \delta^2)}} = 1/\sqrt{1 - \delta}. \end{aligned}$$

which establishes (62)-(63).

To turn these estimates into mathematics we compute for $0 < \tau < 1$, and $\epsilon > 0$

$$\begin{aligned} \mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ &= \mathbb{E} \left\{ \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ \mid \left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| > \epsilon \right\} \mathbb{P} \left(\left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| > \epsilon \right) \\ &\quad + \mathbb{E} \left\{ \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ \mid \left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right\} \mathbb{P} \left(\left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right) \\ \mathbb{P} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} > \tau \right) &= \mathbb{P} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} > \tau \mid \left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| > \epsilon \right) \mathbb{P} \left(\left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| > \epsilon \right) \\ &\quad + \mathbb{P} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} > \tau \mid \left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right) \mathbb{P} \left(\left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right) \end{aligned}$$

For any $0 < \epsilon < \min(1 - \tau, \frac{1}{2})$ we get

$$\begin{aligned} \mathbb{P} \left(\left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right) &\stackrel{(a)}{\geq} 1 - 2e^{-cn\epsilon^2} \\ \mathbb{E} \left\{ \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ \mid \left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right\} &\geq 1 - \tau - \epsilon > 0 \\ \mathbb{P} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} > \tau \mid \left| \frac{\|\mathbf{h}\|_2}{\sqrt{n}} - 1 \right| \leq \epsilon \right) &= 1 \\ \mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ &\geq (1 - \tau - \epsilon)(1 - 2e^{-cn\epsilon^2}) \\ \mathbb{P} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} > \tau \right) &\geq (1 - 2e^{-cn\epsilon^2}) \end{aligned}$$

where (a) follows from Lemma A.1(b) by noting that for $0 < \epsilon \leq \frac{1}{2}$, $\sqrt{\frac{1}{1-\epsilon}} \leq 1 + \epsilon$. Hence, with $\epsilon = (1 - \tau)/n^{1/4}$,

$$F_1(\tau; n) \geq (1 - \tau)^2(1 - n^{-1/4})^2(1 - 2e^{-c(1-\tau)^2n^{1/2}})^2 \quad (77)$$

$$F_2(\tau; n) \geq \tau(1 - \tau)(1 - n^{-1/4})(1 - 2e^{-c(1-\tau)^2n^{1/2}})^2 \quad (78)$$

For an upper bound, denote $c_n \stackrel{\text{def}}{=} \frac{1}{\sqrt{n}} \mathbb{E} \|\mathbf{h}\|_2$. Since $\frac{n}{\sqrt{n+1}} \leq \mathbb{E} \|\mathbf{h}\|_2 \leq \sqrt{n}$ (cf (53) for the lower bound, Jensen's inequality for the upper bound) we have $1 \geq c_n \geq \sqrt{\frac{n}{n+1}} \geq 1 - \frac{1}{2n}$ so that $0 \leq 2(1 - c_n) \leq 1/n$. By Jensen's inequality, for $0 < \tau < 1$,

$$\begin{aligned} \mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+ &\leq \sqrt{\mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)_+^2} \leq \sqrt{\mathbb{E} \left(\frac{\|\mathbf{h}\|_2}{\sqrt{n}} - \tau \right)^2} = \sqrt{1 - 2c_n\tau + \tau^2} \\ &= \sqrt{(1 - \tau)^2 + 2(1 - c_n)\tau} \leq \sqrt{(1 - \tau)^2 + \tau/n} = (1 - \tau)\sqrt{1 + \frac{\tau}{(1-\tau)^2n}} \end{aligned}$$

so that

$$F_1(\tau; n) \leq (1 - \tau)^2 + \tau/n \quad (79)$$

$$F_2(\tau; n) \leq \tau(1 - \tau)\sqrt{1 + \frac{\tau}{(1-\tau)^2 n}} \leq \tau(1 - \tau)\left(1 + \frac{\tau}{(1-\tau)^2 n}\right) \quad (80)$$

Combining all of the above yields (75)-(76).

- **Step 9. (Proof of Lemma B.6).** By (58) and (59) we have for any $n \geq 2/(1 - \delta)$

$$0 < t_{\min}(\delta) \stackrel{\text{def}}{=} \theta^{-1}\left(\frac{1+\delta}{2}\right) \leq t_1^*(\delta; n) \leq 2\theta^{-1}(\delta) \stackrel{\text{def}}{=} t_{\max}(\delta).$$

By the continuity of θ and its strict monotonicity, we have

$$V(\delta) \stackrel{\text{def}}{=} \inf_{t \in [t_{\min}(\delta), t_{\max}(\delta)]} \{\theta(t) - \theta(2t)\} > 0 \quad (81)$$

Applying (69) (convexity of $D_1(t)$) with $h = t = t_1^*$ yields $D_1(2t_1^*) \geq D_1(t_1^*) + t_1^* D_1'(t_1^*)$, hence

$$\begin{aligned} -t_1^* D_1'(t_1^*; n) &\geq D_1(t_1^*; n) - D_1(2t_1^*; n) \stackrel{(70)\&(72)}{\geq} \bar{D}_1(t_1^*; n) - \bar{D}_1(2t_1^*; n) - \frac{1}{n} \\ &\stackrel{(71)}{=} \theta(t_1^*) - \theta(2t_1^*) - \frac{1}{n} \\ &\geq V(\delta) - \frac{1}{n}. \end{aligned}$$

For $n \geq N(\delta) \stackrel{\text{def}}{=} \max(2/V(\delta), 2/(1 - \delta))$ we obtain $-t_1^* D_1'(t_1^*; n) \geq V(\delta)/2 \stackrel{\text{def}}{=} \gamma(\delta) > 0$ which establishes (64).

By (70)-(71)-(72) we have $\theta(t) - 1/n \leq D_1(t; n) \leq \theta(t)$ for all t and n . Hence, the sequence of convex differentiable functions $\{D_1(\cdot; n)\}_n$ converges uniformly to the convex and smooth function $\theta(t)$. It is a classical exercise in convex analysis to show that this implies the convergence of derivatives: indeed, since $D_1(\cdot; n)$ is convex we have for $h_n = n^{-1/2}$ with $n > t^{-2}$ (so that $t - h_n > 0$):

$$\begin{aligned} D_1'(t; n) &\leq \frac{D_1(t + h_n; n) - D_1(t; n)}{h_n} \leq \frac{\theta(t + h_n) - \theta(t) + \frac{1}{n}}{h_n} = \frac{\theta(t + h_n) - \theta(t)}{h_n} + \frac{1}{\sqrt{n}} \\ D_1'(t; n) &\geq \frac{D_1(t; n) - D_1(t - h_n; n)}{h_n} \geq \frac{\theta(t) - \theta(t - h_n) - \frac{1}{n}}{h_n} = \frac{\theta(t) - \theta(t - h_n)}{h_n} - \frac{1}{\sqrt{n}} \end{aligned}$$

It follows that $\lim_{n \rightarrow \infty} D_1'(t; n) = \theta'(t)$. Just as for the case $p = 2$ we have shown that $\lim_{n \rightarrow \infty} \{D_1(t; n) - (t/2)D_1'(t; n)\} = \theta(t) - (t/2)\theta'(t)$. By Lemma A.3, $\theta(t) - (t/2)\theta'(t) = \text{erfc}(t/\sqrt{2})$ so the unique $t = t_1^*(\delta)$ such that $\theta(t) - (t/2)\theta'(t) = \delta$ is

$$t_1^*(\delta) = \sqrt{2}\text{erfc}^{-1}(\delta).$$

With the same reasoning as for the case $p = 2$ above we get that

$$\begin{aligned} \lim_{n \rightarrow \infty} t_1^*(\delta; n) &= t_1^*(\delta) \\ \lim_{n \rightarrow \infty} -t_1^* D_1'(t_1^*; n) &= -\frac{t_1^*(\delta)}{2}\theta'(t_1^*(\delta)) > 0 \\ \lim_{n \rightarrow \infty} D_1(t_1^*; n) &= \theta(t_1^*(\delta)) \\ \lim_{n \rightarrow \infty} \alpha_1^*(\delta; n) &= \sqrt{\frac{\theta(t_1^*(\delta))}{\delta(\delta - \theta(t_1^*(\delta)))}}. \end{aligned}$$

Since $\operatorname{erfc}(t_1^*/\sqrt{2}) = \delta$ we have $\theta(t_1^*) = \delta - \left(\sqrt{\frac{2}{\pi}} e^{-\frac{(t_1^*)^2}{2}} t_1^* - \delta(t_1^*)^2 \right)$ and we conclude that

$$\lim_{n \rightarrow \infty} \alpha_1^*(\delta; n) = \sqrt{\frac{\delta - (\delta - \theta(t_1^*))}{\delta(\delta - \theta(t_1^*))}} = \sqrt{\frac{1}{\sqrt{\frac{2}{\pi}} e^{-\frac{(t_1^*)^2}{2}} t_1^* - \delta(t_1^*)^2} - \frac{1}{\delta}}} \quad (82)$$

□

Lemma B.7 (Deterministic properties of $\kappa(\alpha, \beta)$). *Consider $1 \leq m < n$ two integers, $\delta \stackrel{\text{def}}{=} (m-1)/n$, $1 \leq p \leq \infty$, $\kappa(\alpha, \beta)$ defined in (43), $D_p(t)$ defined in (52). The following hold:*

1. *The function $\kappa(\alpha, \beta)$ is convex-concave and proper on $[0, \infty) \times [0, \infty)$, hence the function*

$$\kappa(\alpha) \stackrel{\text{def}}{=} \sup_{0 \leq \beta \leq 1} \kappa(\alpha, \beta) \quad (83)$$

is convex on $[0, \infty)$, and for any $A > 0$ the function

$$\underline{\kappa}_A(\beta) \stackrel{\text{def}}{=} \inf_{0 \leq \alpha \leq A} \kappa(\alpha, \beta) \quad (84)$$

is concave on $[0, \infty)$.

2. *The scalar $t^* = t_p^*(\delta; n)$ (cf. Lemma B.3-Property 7) is well defined, with $D_p(t^*) < \delta$.*
3. *Define*

$$\alpha^* = \alpha^*(\delta; n) \stackrel{\text{def}}{=} \sqrt{\frac{D_p(t^*)}{\delta - D_p(t^*)}} \quad (85)$$

For $A > \alpha^$ and $\lambda \leq t^*$ we have*

$$\arg \min_{\alpha: 0 \leq \alpha \leq A} \max_{\beta: 0 \leq \beta \leq 1} \kappa(\alpha, \beta) = \alpha^*. \quad (86)$$

The corresponding optimal β is $\beta^ = \beta^*(\lambda, \delta; n) \stackrel{\text{def}}{=} \lambda/t^*$.*

4. *For $A > \alpha^*$, $\lambda \leq t^*$, $0 < \epsilon \leq \max(\alpha^*, A - \alpha^*)$ we have*

$$\inf_{|\alpha - \alpha^*| \geq \epsilon} \kappa(\alpha) - \kappa(\alpha^*) = \inf_{|\alpha - \alpha^*| = \epsilon} \kappa(\alpha) - \kappa(\alpha^*) \geq \omega(\epsilon) = \omega_p(\epsilon; n, \delta, \lambda) \stackrel{\text{def}}{=} \frac{\epsilon^2}{2} \frac{\lambda \delta / t^*}{(1 + \delta(\alpha^* + \epsilon)^2)^{3/2}}. \quad (87)$$

Observe that for the considered range of λ and ϵ , we have $\omega(\epsilon) \leq 1/2$.

Proof of Lemma B.7.

1. It is obvious that κ is proper. One easily checks that $\alpha \mapsto \sqrt{\delta \alpha^2 + 1}$ is convex by checking the non-negativity of its second derivative, hence κ is convex in α . The concavity in β follows from the convexity of $\beta \mapsto \Delta_p(\beta; \mathbf{h}, \lambda)$ which is a distance to a convex set [BV04, Example 3.16], and the fact that the expectation $\Delta_p(\beta, \lambda)$ of a convex function is convex. As a result, $\kappa(\alpha)$ is convex and $\underline{\kappa}(\beta)$ is concave.

2. We have $\delta \stackrel{\text{def}}{=} (m-1)/n < 1$, and by Lemma B.3, Property 4, $D_p(0) \geq \frac{n}{n+1}$. Because we consider the underdetermined case, $1 \leq m < n$, we have $n \geq 2$ and

$$\delta \leq (n-2)/n \leq n/(n+1) \leq D_p(0).$$

3. Since $\kappa(\alpha, \beta)$ is convex-concave and proper, and the constraint sets in (86) are convex and compact, we can change the order of maximization and minimization [Sio58, Corollary 3.3].

- Consider the minimization over α first.

For $\beta = 0$ and any $\lambda \geq 0$ we have $\Delta_p(\beta, \lambda) = 0$ hence $\underline{\kappa}_A(0) = \inf_{0 \leq \alpha \leq A} \kappa(\alpha, 0) = 0$.

For $\beta > 0$, observing that $\Delta_p(\beta, \lambda) = \beta \sqrt{D_p(\lambda/\beta)}$ we rewrite

$$\kappa(\alpha, \beta) = \beta(\sqrt{\delta\alpha^2 + 1} - \alpha\sqrt{D_p(\lambda/\beta)}).$$

With $D_p^{-1}(y) \stackrel{\text{def}}{=} \inf\{t : D_p(t) \leq y\}$ the reciprocal of the strictly decreasing function D_p , we have

$$D_p(\lambda/\beta) \leq \frac{A^2\delta^2}{1 + \delta A^2}$$

if and only if $0 < \beta \leq \tilde{\beta} \stackrel{\text{def}}{=} \frac{\lambda}{D_p^{-1}(A^2\delta^2/(1+\delta A^2))}$. Since $\frac{A^2\delta^2}{1+\delta A^2} < \delta$, this implies that for $0 < \beta \leq \tilde{\beta}$ we can define

$$\tilde{\alpha}(\beta) \stackrel{\text{def}}{=} \sqrt{\frac{D_p(\lambda/\beta)}{\delta(\delta - D_p(\lambda/\beta))}}. \quad (88)$$

and check that $\tilde{\alpha}(\beta) \leq A$. By studying the sign of $\frac{\partial \kappa(\alpha, \beta)}{\partial \alpha} = \beta \left(\frac{\delta\alpha}{\sqrt{\delta\alpha^2 + 1}} - \sqrt{D_p(\lambda/\beta)} \right)$, we get that $\alpha \mapsto \kappa(\alpha, \beta)$ has a unique minimizer on $[0, A]$ which is precisely $\alpha^*(\beta) = \tilde{\alpha}(\beta)$. It follows that for $0 < \beta \leq \tilde{\beta}$ we have:

$$\underline{\kappa}_A(\beta) = \min_{0 \leq \alpha \leq A} \kappa(\alpha, \beta) = \kappa(\alpha^*(\beta), \beta) = \beta \sqrt{\frac{\delta - D_p(\lambda/\beta)}{\delta}}. \quad (89)$$

- Consider now the maximization over β . For $0 < \beta \leq \tilde{\beta}$ the expression of $\underline{\kappa}_A(\beta)$ is given by (89). The sign of $\underline{\kappa}'_A(\beta)$ is that of $\delta - g(\lambda/\beta)$ with $g(t) \stackrel{\text{def}}{=} D_p(t) - \frac{t}{2}D'_p(t)$ as in Lemma B.3-Property 7. Hence, we have: $\underline{\kappa}'_A(\beta) > 0$ if $t = \lambda/\beta > t^*$ (that is to say if $\beta < \beta^* \stackrel{\text{def}}{=} \lambda/t^*$); $\underline{\kappa}'_A(\beta) < 0$ if $\beta > \beta^*$; and $D_p(t^*) < \delta$.

We check that $A > \alpha^*$ implies $D_p(t^*) < A^2\delta^2/(1 + \delta A^2)$ i.e. $\beta^* < \tilde{\beta}$. Combined with the fact that $\underline{\kappa}(0) = 0$, this shows that the supremum of $\underline{\kappa}(\beta)$ over $[0, \tilde{\beta}]$ is indeed achieved uniquely at β^* . This also implies that $\underline{\kappa}(\beta)$ is strictly decreasing for $\beta^* < \beta \leq \tilde{\beta}$. Being concave, $\underline{\kappa}_A(\beta)$ must be also strictly decreasing for $\beta \geq \tilde{\beta}$, so the supremum over $[0, \infty)$ is indeed achieved at β^* . Since $\lambda \leq t^*$ we further have $\beta^* \leq 1$ hence this is also the supremum over $\beta \in [0, 1]$.

To summarize, the optimal β is $\beta^* = \lambda/t^*$, and the corresponding optimal α is given as

$$\alpha^*(\beta^*) = \tilde{\alpha}(\lambda/t^*) = \sqrt{\frac{D_p(t^*)}{\delta(\delta - D_p(t^*))}}. \quad (90)$$

4. The assumption $\epsilon \leq \max(\alpha^*, A - \alpha^*)$ ensures that the set $\{\alpha : |\alpha - \alpha^*| \geq \epsilon, 0 \leq \alpha \leq A\}$ is not empty. Since $\kappa(\alpha)$ is convex on $[0, \infty)$ with its minimum at α^* , we have

$$\omega(\epsilon) \stackrel{\text{def}}{=} \inf_{\substack{\alpha: |\alpha - \alpha^*| \geq \epsilon \\ 0 \leq \alpha \leq A}} \kappa(\alpha) - \kappa(\alpha^*) = \min_{\substack{\alpha: |\alpha - \alpha^*| = \epsilon \\ 0 \leq \alpha \leq A}} \kappa(\alpha) - \kappa(\alpha^*)$$

Since $A > \alpha^*$ and $\lambda \leq t^*$, we have $0 < \beta^* < 1$. The second derivative of $\alpha \mapsto \kappa(\alpha, \beta^*)$ with respect to α reads:

$$\frac{\partial^2 \kappa(\alpha, \beta^*)}{\partial \alpha^2} = \frac{\beta^* \delta}{(1 + \delta \alpha^2)^{3/2}} > 0.$$

This implies that on $[0, A] \cap [\alpha^* - \epsilon, \alpha^* + \epsilon]$ the function $\alpha \mapsto \kappa(\alpha, \beta^*)$ is strongly convex with strong convexity modulus $\frac{\beta^* \delta}{(1 + \delta(\alpha^* + \epsilon)^2)^{3/2}}$. Since $\alpha \mapsto \kappa(\alpha, \beta^*)$ is minimum at α^* , it holds that

$$\kappa(\alpha^* \pm \epsilon, \beta^*) \geq \kappa(\alpha^*, \beta^*) + \frac{\epsilon^2}{2} \frac{\beta^* \delta}{(1 + \delta(\alpha^* + \epsilon)^2)^{3/2}}$$

Furthermore, from the definition of $\kappa(\alpha)$ and β^* , we have that $\kappa(\alpha) \geq \kappa(\alpha, \beta^*)$ for any α , with equality for $\alpha = \alpha^*$. The claim therefore follows.

Since $\lambda \leq t^*$ and $\epsilon \leq \alpha^*$ we have $\omega(\epsilon) \leq \frac{\delta(\alpha^*)^2}{2(1 + \delta(\alpha^*)^2)^{3/2}} = f((1 + \delta(\alpha^*)^2)^{-1/2})$ with $f(u) \stackrel{\text{def}}{=} \frac{1}{2}(1/u^2 - 1)u^3 = \frac{u - u^3}{2} \leq 1/2$ for $0 < u \leq 1$. Hence, $\omega(\epsilon) \leq 1/2$.

□

Invoking a lemma from [HP93] we show that $\arg \min_{\alpha} \phi(\alpha)$ concentrates around $\arg \min_{\alpha} \kappa(\alpha)$.

Lemma B.8 ([HP93, Lemma 2]). *Let $f(t)$ be a random convex function on some open set $S \subset \mathbb{R}^p$, and let t_f be (one of) its minimizer(s). Consider another function $g(t)$ (which we interpret as approximating f), such that it has a unique argmin t_g . Then for each $\epsilon > 0$, we have that:*

$$\mathbb{P}[\|t_f - t_g\| \geq \epsilon] \leq \mathbb{P}\left[\sup_{\|s - t_g\| \leq \epsilon} \|f(s) - g(s)\| \geq \frac{1}{2} \inf_{\|s - t_g\| = \epsilon} (g(s) - g(t_g))\right] \quad (91)$$

The role of $f(t)$ (resp. t_f) will be played by $\phi(\alpha)$ (resp. α_{ϕ}^*), and the role of $g(t)$ (resp. t_g) by $\kappa(\alpha)$ (resp. α^*).

Lemma B.9. *Let $A > \alpha^*$, $\lambda \leq t^*$, with α^* defined as in Lemma B.7-Equation (86) and $t^* = t_p^*(\delta; n)$ as in Lemma B.3-Property 7. Consider the random function $\phi(\alpha)$ defined as in Corollary B.1, and*

$$\alpha_{\phi}^* \stackrel{\text{def}}{=} \arg \min_{0 \leq \alpha \leq A} \phi(\alpha). \quad (92)$$

For $0 < \epsilon \leq \max(\alpha^*, A - \alpha^*)$ we have

$$\mathbb{P}\left[|\alpha_{\phi}^* - \alpha^*| \geq \epsilon\right] \leq \zeta\left(n, \frac{\omega(\epsilon)}{2}\right) \quad (93)$$

with $\zeta(n, \epsilon) = \zeta(n, \epsilon; A, \delta)$ defined in (42) and $\omega(\epsilon)$ defined in (87).

Proof. Since $0 < \epsilon \leq \max(\alpha^*, A - \alpha^*)$ the set $\{\alpha : 0 \leq \alpha \leq A, |\alpha - \alpha^*| \leq \epsilon\}$ is a non-empty subset of $[0, A]$, and

$$\sup_{\alpha: 0 \leq \alpha \leq A, |\alpha - \alpha^*| \leq \epsilon} |\phi(\alpha) - \kappa(\alpha)| \leq \sup_{0 \leq \alpha \leq A} |\phi(\alpha) - \kappa(\alpha)|.$$

Since ϕ is a random convex function, we can apply Lemma B.8 to obtain

$$\begin{aligned} \mathbb{P} [|\alpha_\phi^* - \alpha^*| \geq \epsilon] &\leq \mathbb{P} \left[\sup_{\alpha: 0 \leq \alpha \leq A, |\alpha - \alpha^*| \leq \epsilon} |\phi(\alpha) - \kappa(\alpha)| \geq \frac{\omega(\epsilon)}{2} \right] \leq \mathbb{P} \left[\sup_{0 \leq \alpha \leq A} |\phi(\alpha) - \kappa(\alpha)| \geq \frac{\omega(\epsilon)}{2} \right] \\ &\leq \zeta(n, \omega(\epsilon)/2), \end{aligned}$$

where we used that $\omega(\epsilon) \leq \inf_{\alpha: 0 \leq \alpha \leq A, |\alpha - \alpha^*| = \epsilon} \kappa(\alpha) - \kappa(\alpha^*)$, and the last inequality follows from Corollary B.1 – which we can use since $\omega(\epsilon)/2 \leq 1/4 < 2$. \square

Lemma B.10. *Let $A > \alpha^*$, $\lambda \leq t^*$ with α^* defined as in Lemma B.7-Equation (86) and $t^* = t_p^*(\delta; n)$ as in Lemma B.3-Property 7. For $0 \leq \epsilon \leq \max(\alpha^*, A - \alpha^*)$, consider the optimal cost of the auxiliary optimization (27) with an altered order of minimization and maximization and $\|\mathbf{z}\|_2$ further restricted to be ϵ -away from α^**

$$\phi_\epsilon(\mathbf{g}, \mathbf{h}) \stackrel{\text{def}}{=} \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\substack{\mathbf{z}: \|\mathbf{z}\| \leq A \\ \|\mathbf{z}\| \notin (\alpha^* - \epsilon, \alpha^* + \epsilon)}} \max_{\mathbf{u}: \|\mathbf{u}\| = \beta} \frac{1}{\sqrt{n}} \|\mathbf{z}\| \mathbf{g}^\top \mathbf{u} - \frac{1}{\sqrt{n}} \|\mathbf{u}\| \mathbf{h}^\top \mathbf{z} - \mathbf{u}^\top \mathbf{e}_1 + \frac{\lambda}{\sqrt{n}} \mathbf{w}^\top \mathbf{z}. \quad (94)$$

With $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, $\mathbf{h} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_n)$, $1 \leq m < n$, $\delta = (m - 1)/n$ we have for any $0 < \eta < 2$: with probability at least $1 - \zeta(n, \eta)$,

$$\begin{aligned} \phi_0(\mathbf{g}, \mathbf{h}) &< \kappa(\alpha^*) + \eta, \\ \phi_\epsilon(\mathbf{g}, \mathbf{h}) &> \kappa(\alpha^*) + \omega(\epsilon) - \eta, \end{aligned} \quad (95)$$

with $\zeta(n, \epsilon) = \zeta(n, \epsilon; A, \delta)$ defined in (42) and $\omega(\epsilon)$ defined in (87).

Proof. Since $0 < \epsilon \leq \max(\alpha^*, A - \alpha^*)$ the set $S_\epsilon \stackrel{\text{def}}{=} \{\alpha : 0 \leq \alpha \leq A, |\alpha - \alpha^*| \leq \epsilon\}$ is a non-empty subset of $[0, A]$. Denote $S_\epsilon^- \stackrel{\text{def}}{=} [0, \alpha^* - \epsilon]$ and $S_\epsilon^+ = [\alpha^* + \epsilon, A]$ its two convex components (at most one of them may be empty). Let $\phi_{S_\epsilon^+}(\mathbf{g}, \mathbf{h})$ the value of (94), but with $\|\mathbf{z}\|$ constrained to lie in S_ϵ^+ (by convention, this is $+\infty$ when $S_\epsilon^+ = \emptyset$). Similarly define $\phi_{S_\epsilon^-}(\mathbf{g}, \mathbf{h})$. When S_ϵ^+ is non-empty, since it is convex, we can effect the same simplifications and min-max swaps as in the proof of Lemma 3.1 (from (28a) to (29d)) to arrive at

$$\phi_{S_\epsilon^+}(\mathbf{g}, \mathbf{h}) = \min_{\alpha: \alpha^* + \epsilon \leq \alpha \leq A} \max_{\beta: 0 \leq \beta \leq 1} \phi(\alpha, \beta; \mathbf{g}, \mathbf{h})$$

and similarly with S_ϵ^- we get when it is non-empty that

$$\phi_{S_\epsilon^-}(\mathbf{g}, \mathbf{h}) = \min_{\alpha: 0 \leq \alpha \leq \alpha^* - \epsilon} \max_{\beta: 0 \leq \beta \leq 1} \phi(\alpha, \beta; \mathbf{g}, \mathbf{h})$$

This shows that $\phi_\epsilon(\mathbf{g}, \mathbf{h}) = \min(\phi_{S_\epsilon^-}(\mathbf{g}, \mathbf{h}), \phi_{S_\epsilon^+}(\mathbf{g}, \mathbf{h})) = \phi_{S_\epsilon}(\mathbf{g}, \mathbf{h})$ where the notation $\phi_{S_\epsilon}(\mathbf{g}, \mathbf{h})$ matches that used in Corollary B.1. Moreover by definition (see Lemma B.7) we have

$$\kappa_{S_\epsilon} \stackrel{\text{def}}{=} \min_{\alpha \in S_\epsilon} \max_{0 \leq \beta \leq 1} \kappa(\alpha, \beta) = \kappa(\alpha^*) + \omega(\epsilon).$$

By Corollary B.1 we have, for $0 < \eta < 2$, with probability at least $1 - \zeta(n, \eta)$: for all $S \subset [0, A]$, $|\phi_S(\mathbf{g}, \mathbf{h}) - \kappa_S| < \eta$. Specializing to $S = [0, A]$ and $S = S_\epsilon$ and combining the above yields

$$\begin{aligned}\phi_0(\mathbf{g}, \mathbf{h}) &= \phi_{[0, A]}(\mathbf{g}, \mathbf{h}) < \kappa_{[0, A]} + \eta = \kappa(\alpha^*) + \eta \\ \phi_\epsilon(\mathbf{g}, \mathbf{h}) &= \phi_{S_\epsilon}(\mathbf{g}, \mathbf{h}) > \kappa_{S_\epsilon} - \eta = \kappa(\alpha^*) + \omega(\epsilon) - \eta.\end{aligned}$$

□

Lemma B.11. *Let $K > \alpha^*/\sqrt{n}$, $\lambda \leq t^*$ with α^* defined as in Lemma B.7-Equation (86) and $t^* = t_p^*(\delta; n)$ as in Lemma B.3-Property 7. Denote by $\tilde{\mathbf{x}}_K$ any optimal solution of (24). For $0 < \epsilon \leq \max(\alpha^*, K\sqrt{n} - \alpha^*)$ we have*

$$\mathbb{P}\left[\left|\|\tilde{\mathbf{x}}_K\| - \frac{\alpha^*}{\sqrt{n}}\right| \geq \frac{\epsilon}{\sqrt{n}}\right] \leq 4\zeta\left(n, \frac{\omega(\epsilon)}{2}\right).$$

with $\zeta(n, \epsilon) = \zeta(n, \epsilon; K\sqrt{n}, \delta)$ defined in (42) and $\omega(\epsilon)$ defined in (87).

Proof. Denote Φ the optimal cost of (24) and Φ_ϵ the corresponding cost when $\mathbf{z} = \mathbf{x}\sqrt{n}$ is further restricted to $\mathcal{S}_\epsilon \stackrel{\text{def}}{=} \{\mathbf{z} : \|\mathbf{z}\|_2 \leq A \text{ and } \left|\|\mathbf{z}\|_2 - \alpha^*\right| \geq \epsilon\}$, with $A \stackrel{\text{def}}{=} K\sqrt{n}$:

$$\Phi_\epsilon = \min_{\mathbf{x}: \|\mathbf{x}\| \leq K, \left|\sqrt{n}\|\mathbf{x}\| - \alpha^*\right| \geq \epsilon} \max_{\mathbf{u}: \|\mathbf{u}\| \leq 1} \mathbf{u}^\top \mathbf{A}\mathbf{x} + \lambda \|\mathbf{x}\|_p - \mathbf{u}^\top \mathbf{e}_1. \quad (96)$$

We now want to show that for $\epsilon > 0$ we have with high probability $\Phi_\epsilon > \Phi = \Phi_0$, because this is equivalent to $\tilde{\mathbf{z}}_K = \tilde{\mathbf{x}}_K\sqrt{n} \in \mathcal{S}'_\epsilon \stackrel{\text{def}}{=} \{\mathbf{z} : \|\mathbf{z}\|_2 \leq A \text{ and } \left|\|\mathbf{z}\|_2 - \alpha^*\right| < \epsilon\}$.

By CGMT part 1 and part 2 (see Theorem A.1 in Appendix A.3), denoting ϕ_ϵ^P (resp. ϕ_ϵ^D) the optimum value of the “primal” (resp. “dual”) auxiliary optimization problem associated to the principal optimization problem (96), we have for any $c \in \mathbb{R}$

$$\mathbb{P}[\Phi_\epsilon < c] \leq 2\mathbb{P}[\phi_\epsilon^P \leq c] \quad \text{and} \quad \mathbb{P}[\Phi_0 > c] \leq 2\mathbb{P}[\phi_0^D \geq c]. \quad (97)$$

For the second one, we use additionally that $\phi_0^D = \phi_0^P$ since we optimize over convex sets (being a ball, \mathcal{S}_0 is convex) and the penalty $\psi(\mathbf{x}, \mathbf{u}) = \lambda \|\mathbf{x}\|_p - \mathbf{u}^\top \mathbf{e}_1$ is convex-concave (see e.g. [Sio58, Corollary 3.3]).

Let $C(\mathbf{z}, \mathbf{u}, \mathbf{w}) = C(\mathbf{z}, \mathbf{u}, \mathbf{z}; \mathbf{g}, \mathbf{h})$ be the objective function in (27) and (94), so that (94) becomes (with $A \stackrel{\text{def}}{=} K\sqrt{n}$)

$$\phi_\epsilon = \max_{\substack{\beta: 0 \leq \beta \leq 1 \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\substack{\mathbf{z}: \|\mathbf{z}\| \leq A \\ \mathbf{z} \in \mathcal{S}_\epsilon}} \max_{\mathbf{u}: \|\mathbf{u}\| = \beta} C(\mathbf{z}, \mathbf{u}, \mathbf{w}),$$

and the optimal cost of (27) reads $\phi = \phi_0$. With these notations, we have (noting that $\max \min \leq \min \max$ is always true)

$$\begin{aligned}\phi_\epsilon^P &\stackrel{\text{def}}{=} \min_{\substack{\mathbf{z}: \|\mathbf{z}\| \leq A \\ \mathbf{z} \in \mathcal{S}_\epsilon}} \max_{\substack{\mathbf{u}: \|\mathbf{u}\| \leq 1, \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} C(\mathbf{z}, \mathbf{u}, \mathbf{w}) = \min_{\substack{\mathbf{z}: \|\mathbf{z}\| \leq A \\ \mathbf{z} \in \mathcal{S}_\epsilon}} \max_{\beta: 0 \leq \beta \leq 1, \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1} \max_{\mathbf{u}: \|\mathbf{u}\|_2 = \beta} C(\mathbf{z}, \mathbf{u}, \mathbf{w}) \\ &\geq \max_{\substack{\beta: 0 \leq \beta \leq 1, \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\substack{\mathbf{z}: \|\mathbf{z}\| \leq A \\ \mathbf{z} \in \mathcal{S}_\epsilon}} \max_{\mathbf{u}: \|\mathbf{u}\|_2 = \beta} C(\mathbf{z}, \mathbf{u}, \mathbf{w}) = \phi_\epsilon\end{aligned} \quad (98)$$

and

$$\begin{aligned} \phi_0^D &\stackrel{\text{def}}{=} \max_{\substack{\mathbf{u}: \|\mathbf{u}\| \leq 1, \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\mathbf{z}: \|\mathbf{z}\| \leq A} C(\mathbf{z}, \mathbf{u}, \mathbf{w}) = \max_{\substack{\beta: 0 \leq \beta \leq 1, \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \max_{\mathbf{u}: \|\mathbf{u}\|_2 = \beta} \min_{\|\mathbf{z}\| \leq A} C(\mathbf{z}, \mathbf{u}, \mathbf{w}) \\ &\leq \max_{\substack{\beta: 0 \leq \beta \leq 1, \\ \mathbf{w}: \|\mathbf{w}\|_{p^*} \leq 1}} \min_{\mathbf{z}: \|\mathbf{z}\| \leq A} \max_{\mathbf{u}: \|\mathbf{u}\|_2 = \beta} C(\mathbf{z}, \mathbf{u}, \mathbf{w}) = \phi_0. \end{aligned} \quad (99)$$

Denote $\underline{\phi}_\epsilon \stackrel{\text{def}}{=} \kappa(\alpha^*) + \omega(\epsilon)$ and $\bar{\phi} \stackrel{\text{def}}{=} \kappa(\alpha^*)$ and use the above with $c_1 = \underline{\phi}_\epsilon - \eta$ and $c_2 = \bar{\phi} + \eta$ where $\eta > 0$ is arbitrary to get

$$\begin{aligned} \mathbb{P}[\Phi_\epsilon < \underline{\phi}_\epsilon - \eta] &\leq 2\mathbb{P}[\phi_\epsilon^P \leq \underline{\phi}_\epsilon - \eta] \leq 2\mathbb{P}[\phi_\epsilon \leq \underline{\phi}_\epsilon - \eta], \\ \mathbb{P}[\Phi > \bar{\phi} + \eta] &\leq 2\mathbb{P}[\phi^D \geq \bar{\phi} + \eta] \leq 2\mathbb{P}[\phi_0 \geq \bar{\phi} + \eta]. \end{aligned} \quad (100)$$

Consider the event $\mathcal{E} = \{\Phi_\epsilon \geq \underline{\phi}_\epsilon - \eta \text{ and } \Phi \leq \bar{\phi} + \eta\}$. For $0 < \eta < (\underline{\phi}_\epsilon - \bar{\phi})/2 = \omega(\epsilon)/2$ we have $c_1 > c_2$ hence this event implies that $\tilde{\mathbf{z}}_K \in \mathcal{S}'_\epsilon$, which is what we wanted to prove. For such η , since $\omega(\epsilon)/2 \leq 1/4 < 2$ we can use Lemma B.10 and a union bound to obtain that this event happens with probability at least $1 - 4\zeta(n, \eta)$. Hence, for any $0 < \eta < \omega(\epsilon)/2$ we have $\mathbb{P}(\tilde{\mathbf{z}}_K \notin \mathcal{S}_\epsilon) \leq 4\zeta(n, \eta)$. By continuity of $\eta \mapsto \zeta(n, \eta)$ we take the limit when η tends to $\omega(\epsilon)/2$. \square

Lemma B.12. *Let $K > \alpha^*/\sqrt{n}$, $\lambda \leq t^*$ with α^* defined as in Lemma B.7-Equation (86) and $t^* = t_p^*(\delta; n)$ as in Lemma B.3-Property 7. Denote by $\tilde{\mathbf{x}}_K$ any optimal solution of the random bounded problem (24) and $\tilde{\mathbf{x}}$ any optimal solution of the random unbounded problem (22). For $0 < \epsilon \leq \min(\alpha^*, K\sqrt{n} - \alpha^*)$ (NB: here the upper bound on ϵ is the min, not the max) we have*

$$\mathbb{P}[\tilde{\mathbf{x}}_K \neq \tilde{\mathbf{x}}] \leq 4\zeta\left(n, \frac{\omega(\epsilon)}{2}\right),$$

and

$$\mathbb{P}\left[\left|\|\tilde{\mathbf{x}}\| - \frac{\alpha^*}{\sqrt{n}}\right| \geq \frac{\epsilon}{\sqrt{n}}\right] \leq 4\zeta\left(n, \frac{\omega(\epsilon)}{2}\right).$$

with $\zeta(n, \epsilon) = \zeta(n, \epsilon; K\sqrt{n}, \delta)$ defined in (42) and $\omega(\epsilon)$ defined in (87).

Proof. To handle the case of non-unique solutions, $\tilde{\mathbf{x}}$ (resp. $\tilde{\mathbf{x}}_K$) may denote the convex set of solutions of the respective convex optimization problems. The property $\tilde{\mathbf{x}} \neq \tilde{\mathbf{x}}_K$ then means that the sets do not intersect, and inequalities such as $f(\tilde{\mathbf{x}}) > c$ are meant to hold for all elements of the set $\tilde{\mathbf{x}}$.

We first prove, by contradiction, that if $\tilde{\mathbf{x}} \neq \tilde{\mathbf{x}}_K$, then necessarily $\|\tilde{\mathbf{x}}_K\| = K$. Suppose the opposite: despite the fact that $\tilde{\mathbf{x}} \neq \tilde{\mathbf{x}}_K$, we have $\|\tilde{\mathbf{x}}_K\| < K$. Since $\tilde{\mathbf{x}} \neq \tilde{\mathbf{x}}_K$ we have $\|\tilde{\mathbf{x}}\| > K$. Denoting the lasso objective in (22) by $\theta(\mathbf{x})$, this means that $\theta(\tilde{\mathbf{x}}) < \theta(\tilde{\mathbf{x}}_K)$. By convexity of θ it follows that all points on the line segment $\mathbf{x}_\nu = \nu\tilde{\mathbf{x}}_K + (1 - \nu)\tilde{\mathbf{x}}$, $0 \leq \nu \leq 1$ satisfy

$$\theta(\mathbf{x}_\nu) \leq \theta(\tilde{\mathbf{x}}_K). \quad (101)$$

Since $\|\mathbf{x}_0\| > K$ and $\|\mathbf{x}_1\| < K$, by continuity there exists $\nu \in (0, 1)$ such that $\|\mathbf{x}_\nu\| = K$. Further, by (101), \mathbf{x}_ν is optimizing the bounded problem (24), contradicting our assumption.

The contraposition of what we just established is that, if $\|\tilde{\mathbf{x}}_K\| < K$ then $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_K$. In particular, since we assume $K\sqrt{n} > \alpha^* + \epsilon$, we have: if $|\sqrt{n}\|\tilde{\mathbf{x}}_K\| - \alpha^*| < \epsilon$ then $\|\tilde{\mathbf{x}}_K\| <$

K hence $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_K$ and $|\sqrt{n}\|\tilde{\mathbf{x}}\| - \alpha^*| = |\sqrt{n}\|\tilde{\mathbf{x}}_K\| - \alpha^*| < \epsilon$. It follows that $\mathbb{P}[\tilde{\mathbf{x}}_K \neq \tilde{\mathbf{x}}] \leq \mathbb{P}[|\sqrt{n}\|\tilde{\mathbf{x}}_K\| - \alpha^*| \geq \epsilon]$ and

$$\mathbb{P}[|\sqrt{n}\|\tilde{\mathbf{x}}\| - \alpha^*| \geq \epsilon] \leq \mathbb{P}[|\sqrt{n}\|\tilde{\mathbf{x}}_K\| - \alpha^*| \geq \epsilon].$$

We conclude using Lemma B.11. \square

Lemma B.13. *With Δ_p defined as in Lemma B.2, the function $f_{\mathbf{h}}(\beta) \stackrel{\text{def}}{=} |\Delta_p(\beta; \mathbf{h}, \lambda) - \Delta_p(\beta; n, \lambda)|$ is $\max\{\|\mathbf{h}\|/\sqrt{n}, 1\}$ -Lipschitz in β .*

Proof. We omit the dependency in λ for brevity and write by definition

$$|\Delta_p(\beta_1; \mathbf{h}) - \Delta_p(\beta_2; \mathbf{h})| = \frac{1}{\sqrt{n}} |\text{dist}(\beta_1 \mathbf{h}, \mathcal{C}) - \text{dist}(\beta_2 \mathbf{h}, \mathcal{C})| \leq \frac{1}{\sqrt{n}} \|\beta_1 \mathbf{h} - \beta_2 \mathbf{h}\| \leq \frac{\|\mathbf{h}\|}{\sqrt{n}} |\beta_1 - \beta_2|,$$

where we used the fact that the Euclidean distance to a convex set is 1-Lipschitz with respect to the Euclidean metric. Further, we have

$$\begin{aligned} |\Delta_p(\beta_1) - \Delta_p(\beta_2)| &= |\mathbb{E}(\Delta_p(\beta_1; \mathbf{h}) - \Delta_p(\beta_2; \mathbf{h}))| \leq \mathbb{E} |(\Delta_p(\beta_1; \mathbf{h}) - \Delta_p(\beta_2; \mathbf{h}))| \\ &\leq \frac{|\beta_1 - \beta_2|}{\sqrt{n}} \mathbb{E} \|\mathbf{h}\| \leq |\beta_1 - \beta_2|. \end{aligned}$$

To conclude, note that the Lipschitz constant of the difference of two Lipschitz functions does not exceed the largest of the two Lipschitz constants, and that taking the absolute value does not change it. \square

References

- [Bar05] A. Barvinok, “Math 710: Measure concentration,” *Lecture notes*, 2005.
- [BV04] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [CDS03] M. Chiani, D. Dardari, and M. K. Simon, “New exponential bounds and approximations for the computation of error probability in fading channels,” *IEEE Transactions on Wireless Communications*, vol. 24, no. 5, pp. 840–845, 2003.
- [CHKK] P. G. Casazza, A. Heinecke, F. Krahmer, and G. Kutyniok, “Optimally Sparse Frames,” *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7279–7287.
- [DKV13] I. Dokmanić, M. Kolundžija, and M. Vetterli, “Beyond Moore-Penrose: Sparse pseudoinverse,” in *IEEE ICASSP*. IEEE, 2013, pp. 6526–6530.
- [DMA97] G. Davis, S. Mallat, and M. Avellaneda, “Adaptive Greedy Approximations,” *Constr. Approx.*, vol. 13, no. 1, pp. 57–98, 1997.
- [Don06] D. L. Donoho, “Compressed Sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [Dos12] C. Dossal, “A necessary and sufficient condition for exact sparse recovery by minimization,” *Comptes Rendus Mathématique*, vol. 350, no. 1-2, pp. 117–120, Jan. 2012.

- [Ela10] M. Elad, *Sparse and Redundant Representations*, ser. From Theory to Applications in Signal and Image Processing. Springer, 2010.
- [Fuc04] J. J. Fuchs, “On Sparse Representations in Arbitrary Redundant Bases,” *IEEE Trans. Inf. Theory*, vol. 50, no. 6, pp. 1341–1344, Jun. 2004.
- [HMT11] N. Halko, P. G. Martinsson, and J. A. Tropp, “Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions,” *SIAM Rev.*, vol. 53, no. 2, pp. 217–288, May 2011.
- [HP93] N. L. Hjort and D. Pollard, *Asymptotics for Minimisers of Convex Processes*, 1993.
- [KKL12] F. Krahmer, G. Kutyniok, and J. Lemvig, “Sparsity and Spectral Properties of Dual Frames,” *Linear Algebra Appl.*, vol. 439, no. 4, pp. 1–17, Dec. 2012.
- [Led99] M. Ledoux, “Concentration of measure and logarithmic Sobolev inequalities,” in *Séminaire de Probabilités XXXIII*. Berlin, Heidelberg: Springer, Berlin, Heidelberg, 1999, pp. 120–216.
- [LLM13] S. Li, Y. Liu, and T. Mi, “Sparse Dual Frames and Dual Gabor Functions of Minimal Time and Frequency Supports,” *J. Fourier Anal. Appl.*, vol. 19, no. 1, pp. 48–76, 2013.
- [Nat95] B. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM J. Computing*, vol. 25, no. 2, pp. 227–234, 1995.
- [OTH13] S. Oymak, C. Thrampoulidis, and B. Hassibi, “The Squared-Error of Generalized LASSO: A Precise Analysis,” *arXiv*, Nov. 2013.
- [PHSB14] N. Perraudin, N. Holighaus, P. L. Søndergaard, and P. Balazs, “Designing Gabor Windows Using Convex Optimization,” *arXiv*, Jan. 2014.
- [Rud08] M. Rudelson, “Invertibility of random matrices: norm of the inverse,” *Annals of Mathematics*, vol. 168, no. 2, pp. 575–600, 2008.
- [Sio58] M. Sion, “On general minimax theorems,” *Pacific Journal of Mathematics*, vol. 8, no. 1, pp. 171–176, Mar. 1958.
- [TAH16] C. Thrampoulidis, E. Abbasi, and B. Hassibi, “Precise Error Analysis of Regularized M-estimators in High-dimensions,” *arXiv*, Jan. 2016.
- [Tib13] R. J. Tibshirani, “The lasso problem and uniqueness,” *Electronic Journal of Statistics*, vol. 7, no. 0, pp. 1456–1490, 2013.
- [TO15] C. Thrampoulidis and S. Oymak, “Regularized linear regression: A precise analysis of the estimation error,” in *Proceedings of The 28th . . .*, 2015.
- [Ver09] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” in *Compressed Sensing*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge: Cambridge University Press, 2009, pp. 210–268.
- [Ver14] —, “Invertibility of symmetric random matrices,” *Random Structures & Algorithms*, vol. 44, no. 2, pp. 135–182, Mar. 2014.

- [Wai15] M. Wainwright, *Basic tail and concentration bounds, Draft*, 2015. [Online]. Available: https://www.stat.berkeley.edu/~mjwain/stat210b/Chap2.TailBounds_Jan22_2015.pdf
- [Wen48] J. G. Wendel, "Note on the Gamma Function," *Am. Math. Mon.*, vol. 55, no. 9, p. 563, Nov. 1948.