

Chapter 2

Convex Recovery of a Structured Signal from Independent Random Linear Measurements

Joel A. Tropp

Abstract This chapter develops a theoretical analysis of the convex programming method for recovering a structured signal from independent random linear measurements. This technique delivers bounds for the sampling complexity that are similar to recent results for standard Gaussian measurements, but the argument applies to a much wider class of measurement ensembles. To demonstrate the power of this approach, the chapter presents a short analysis of phase retrieval by trace-norm minimization. The key technical tool is a framework, due to Mendelson and coauthors, for bounding a nonnegative empirical process.

2.1 Motivation

Signal reconstruction from random measurements is a central preoccupation in contemporary signal processing. In this problem, we acquire linear measurements of an unknown, structured signal through a random sampling process. Given these random measurements, a standard method for recovering the unknown signal is to solve a convex optimization problem that enforces our prior knowledge about the structure. The basic question is how many measurements suffice to resolve a particular type of structure.

Recent research has led to a comprehensive answer when the measurement operator follows the standard Gaussian distribution [1, 6, 10, 22, 24–26, 29, 31, 33]. The literature also contains satisfying answers for sub-Gaussian measurements [22] and subexponential measurements [18]. Other types of measurement systems are quite common, but we are not aware of a simple approach that allows us to analyze general measurements in a unified way.

This chapter describes an approach that addresses a wide class of convex signal reconstruction problems involving random sampling. To understand these questions, the core challenge is to produce a lower bound on a nonnegative empirical process.

J.A. Tropp (✉)

Department of Computing and Mathematical Sciences, California Institute of Technology, 1200 E. California Blvd. Pasadena, CA, 91125-5000

For this purpose, we rely on a powerful framework, called the *Small Ball Method*, that was developed by Shahar Mendelson and coauthors in a sequence of papers, including [14, 16, 19–21].

To complete the estimates required by Mendelson’s Small Ball Method, we propose a technique based on conic duality. One advantage of this approach is that we can exploit the same insights and calculations that have served so well in the Gaussian setting. We refer to this little argument as the *bowling scheme* in honor of David Gross’s *golfing scheme* [13]. We anticipate that it will offer researchers an effective way to analyze many signal recovery problems with random measurements.

2.1.1 Roadmap

The first half of the chapter summarizes the established analysis of convex signal reconstruction with a Gaussian sampling model. In Section 2.2, we introduce a convex optimization framework for solving structured signal recovery problems with linear measurements, and we present a geometric formulation of the optimality conditions. Section 2.3 specializes to the case where the measurements come from a Gaussian model, and we explain how classical results for Gaussian processes lead to a sharp bound for the number of Gaussian measurements that suffice. These results are framed in terms of a geometric parameter, the conic Gaussian width, associated with the convex optimization problem. Section 2.4 explains how to use duality to obtain a numerically sharp bound for the conic Gaussian width, and it develops two important examples in detail.

In the second half of the chapter, we consider more general sampling models. Section 2.5 introduces Mendelson’s Small Ball Method and the technical arguments that support it. As a first application, in Section 2.6, we use this strategy to analyze signal reconstruction from sub-Gaussian measurements. Section 2.7 presents the bowling scheme, which merges the conic duality estimates with Mendelson’s Small Ball Method. This technique allows us to study more general types of random measurements. Finally, in Section 2.8, we demonstrate the vigor of these ideas by applying them to the phase retrieval problem.

2.2 Signal reconstruction from linear measurements

We begin with a framework that describes many convex optimization methods for recovering a structured signal from linear measurements. Examples include the ℓ_1 minimization approach for identifying a sparse vector and the Schatten 1-norm minimization approach for identifying a low-rank matrix. We develop a simple error bound for convex signal reconstruction by exploiting the geometric formulation of the optimality conditions. This analysis leads us to study the minimum conic singular value of a matrix.

2.2.1 Linear acquisition of data

Let $\mathbf{x}^\natural \in \mathbb{R}^d$ be an unknown but “structured” signal. Suppose that we observe a vector $\mathbf{y}$ in $\mathbb{R}^m$ that consists of m linear measurements of the unknown:

$$\mathbf{y} = \Phi \mathbf{x}^\natural + \mathbf{e}. \quad (2.1)$$

We assume that Φ is a known $m \times d$ sampling matrix, and $\mathbf{e} \in \mathbb{R}^m$ is a vector of unknown errors. Expression (2.1) offers a model for data acquisition that describes a wide range of problems in signal processing, statistics, and machine learning. Our goal is to compute an approximation of the unknown $\mathbf{x}^\natural$ by exploiting our prior knowledge about its structure.

2.2.2 Reconstruction via convex optimization

Convex optimization is a popular approach for recovering a structured vector from linear measurements. Let $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ be a proper convex function¹ that reflects the “complexity” of a signal. Then we can frame the convex program

$$\underset{\mathbf{x} \in \mathbb{R}^d}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad \|\Phi \mathbf{x} - \mathbf{y}\| \leq \eta \quad (2.2)$$

where $\|\cdot\|$ denotes the Euclidean norm and η is a specified bound on the norm of the error $\mathbf{e}$. In words, the optimization problem (2.2) searches for the most structured signal $\mathbf{x}$ that is consistent with the observed data $\mathbf{y}$. In practice, it is common to consider the Lagrangian formulation of (2.2) or to consider a problem where the objective and constraint are interchanged. We can often solve (2.2) and its variants efficiently using standard algorithms.

Remark 2.1 (Alternative programs). The optimization problem (2.2) is not the only type of convex method for signal reconstruction. Suppose that $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ is a gauge, i.e., a function that is nonnegative, positively homogeneous, and convex. Then we may consider the convex program

$$\underset{\mathbf{x} \in \mathbb{R}^d}{\text{minimize}} \quad f(\mathbf{x}) \quad \text{subject to} \quad f^\circ(\Phi^\dagger(\Phi \mathbf{x} - \mathbf{y})) \leq \eta,$$

where f° denotes the polar of the gauge [28, Chap. 15] and † denotes transposition. This reconstruction method submits to an analysis similar to the approach in this note. For example, see [4, Thm. 1].

¹The extended real numbers $\overline{\mathbb{R}} := \mathbb{R} \cup \{\pm\infty\}$. A *proper* convex function takes at least one finite value but never the value $-\infty$.

2.2.3 Examples

Before we continue, let us mention a few structures that arise in applications and the complexity measures that are typically associated with these structures.

Example 2.2 (Sparse vectors). A vector $\mathbf{x}^\natural \in \mathbb{R}^d$ is *sparse* when many or most of its entries are equal to zero. We can promote sparsity by minimizing the ℓ_1 norm $\|\cdot\|_{\ell_1}$. This heuristic leads to a problem of the form

$$\underset{\mathbf{x} \in \mathbb{R}^d}{\text{minimize}} \quad \|\mathbf{x}\|_{\ell_1} \quad \text{subject to} \quad \|\Phi \mathbf{x} - \mathbf{y}\| \leq \eta. \quad (2.3)$$

Sparsity has become a dominant modeling tool in statistics, machine learning, and signal processing.

Example 2.3 (Low-rank matrices). We say that a matrix $\mathbf{X}^\natural \in \mathbb{R}^{d_1 \times d_2}$ has *low rank* when its rank is small compared with minimum of d_1 and d_2 . Suppose that we have acquired noisy measurements

$$\mathbf{y} = \Phi(\mathbf{X}^\natural) + \mathbf{e}, \quad (2.4)$$

where Φ is a linear operator that maps a matrix in $\mathbb{R}^{d_1 \times d_2}$ to a vector in $\mathbb{R}^m$. To reconstruct the unknown low-rank matrix $\mathbf{X}^\natural$, we can minimize the Schatten 1-norm $\|\cdot\|_{S_1}$, which returns the sum of the singular values of a matrix. This heuristic suggests that we consider an optimization problem of the form

$$\underset{\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}}{\text{minimize}} \quad \|\mathbf{X}\|_{S_1} \quad \text{subject to} \quad \|\Phi(\mathbf{X}) - \mathbf{y}\| \leq \eta. \quad (2.5)$$

In recent years, this approach to fitting low-rank matrices has become common.

It is possible to consider many other types of structures. For instance, see [6, 10].

2.2.4 A deterministic error bound for convex recovery

We can obtain a deterministic error bound for the convex reconstruction method (2.2) using a standard geometric analysis. Recall that a *cone* is a set $K \subset \mathbb{R}^d$ that is positively homogeneous: $K = \tau K$ for all $\tau > 0$. A *convex cone* is a cone that is also a convex set. Let us introduce the cone of descent directions of a convex function.

Definition 2.4 (Descent cone). Let $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ be a proper convex function. The *descent cone* $\mathcal{D}(f, \mathbf{x})$ of the function f at a point $\mathbf{x} \in \mathbb{R}^d$ is defined as

$$\mathcal{D}(f, \mathbf{x}) := \bigcup_{\tau > 0} \{\mathbf{u} \in \mathbb{R}^d : f(\mathbf{x} + \tau \mathbf{u}) \leq f(\mathbf{x})\}.$$

The descent cone of a convex function is always a convex cone, but it may not be closed.

We are interested in the behavior of the measurement matrix Φ when it is restricted to a descent cone.

Definition 2.5 (Minimum conic singular value). Let Φ be an $m \times d$ matrix, and let K be a cone in $\mathbb{R}^d$. The minimum singular value of Φ with respect to the cone K is defined as

$$\lambda_{\min}(\Phi; K) := \inf \{ \|\Phi \mathbf{u}\| : \mathbf{u} \in K \cap \mathbf{S}^{d-1} \}$$

where $\mathbf{S}^{d-1}$ is the Euclidean unit sphere in $\mathbb{R}^d$.

The terminology originates in the fact that $\lambda_{\min}(\Phi; \mathbb{R}^d)$ coincides with the usual minimum singular value.

With these definitions at hand, we reach the following basic result.

Proposition 2.6 (A deterministic error bound for convex recovery). Let $\mathbf{x}^\natural$ be a signal in $\mathbb{R}^d$, let Φ be an $m \times d$ measurement matrix, and let $\mathbf{y} = \Phi \mathbf{x}^\natural + \mathbf{e}$ be a vector of measurements in $\mathbb{R}^m$. Assume that $\|\mathbf{e}\| \leq \eta$, and let $\hat{\mathbf{x}}_\eta$ be any solution to the optimization problem (2.2). Then

$$\|\hat{\mathbf{x}}_\eta - \mathbf{x}^\natural\| \leq \frac{2\eta}{\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{x}^\natural))}.$$

This statement is adapted from [6]. For completeness, we include the short proof.

Proof. It is natural to write the decision variable $\mathbf{x}$ in the convex program (2.2) relative to the true unknown: $\mathbf{u} := \mathbf{x} - \mathbf{x}^\natural$. Using expression (2.1) for the measurement vector $\mathbf{y}$, we obtain the equivalent problem

$$\underset{\mathbf{u} \in \mathbb{R}^d}{\text{minimize}} \quad f(\mathbf{x}^\natural + \mathbf{u}) \quad \text{subject to} \quad \|\Phi \mathbf{u} - \mathbf{e}\| \leq \eta. \quad (2.6)$$

Owing to the bound $\|\mathbf{e}\| \leq \eta$, the point $\mathbf{u} = \mathbf{0}$ is feasible for (2.6). Therefore, each optimal point $\hat{\mathbf{u}}$ verifies $f(\mathbf{x}^\natural + \hat{\mathbf{u}}) \leq f(\mathbf{x}^\natural)$. In summary, any optimal point of (2.6) satisfies two conditions:

$$\hat{\mathbf{u}} \in \mathcal{D}(f, \mathbf{x}^\natural) \quad \text{and} \quad \|\Phi \hat{\mathbf{u}} - \mathbf{e}\| \leq \eta.$$

As a consequence, we simply need to determine how far we can travel in a descent direction before we violate the bound constraint. See Figure 2.1 for an illustration of the geometry.

To complete the argument, assume that $\mathbf{u}$ is a nonzero point in $\mathcal{D}(f, \mathbf{x}^\natural)$ that is feasible for (2.6). Then

$$\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{x}^\natural)) \leq \frac{\|\Phi \mathbf{u}\|}{\|\mathbf{u}\|} \leq \frac{\|\Phi \mathbf{u} - \mathbf{e}\| + \|\mathbf{e}\|}{\|\mathbf{u}\|} \leq \frac{2\eta}{\|\mathbf{u}\|}.$$




Fig. 2.1 [Geometry of convex recovery] This diagram illustrates the geometry of the optimization problem (2.6). The cone $\mathcal{D}(f, \mathbf{x}^h)$ contains the directions $\mathbf{u}$ in which f is decreasing at $\mathbf{x}^h$. Assuming that $\|\mathbf{e}\| \leq \eta$, the diagonal tube contains every point $\mathbf{u}$ that satisfies the bound constraint $\|\Phi \mathbf{u} + \mathbf{e}\| \leq \eta$. Each optimal point $\hat{\mathbf{u}}$ for (2.6) lies in the intersection of the tube and the cone.

The first inequality follows from Definition 2.5 of the conic singular value. The second relation is the triangle inequality. The last bound holds because $\mathbf{u}$ satisfies the constraint in (2.6), and we have assumed that $\|\mathbf{e}\| \leq \eta$. Finally, rearrange the display, and rewrite $\mathbf{u}$ in terms of the original decision variable $\mathbf{x}$. $\square$

Although Proposition 2.6 is elegant, it can be difficult to apply because we must calculate the minimum conic singular value of a matrix Φ with respect to a descent cone. This challenge becomes less severe, however, when the matrix Φ is drawn at random.

2.3 A universal error bound for Gaussian measurements

We will study the prospects for convex recovery when the sampling matrix Φ is chosen at random. This modeling assumption arises in signal processing applications where the matrix describes a data-acquisition system that can extract random measurements. This kind of model also appears in statistics and machine learning when each row of the matrix tabulates measured variables for an individual subject in an experiment.

2.3.1 Standard Gaussian measurements

In this section, we treat one of the simplest mathematical models for the $m \times d$ random measurement matrix Φ . We assume that each of the m rows of Φ is drawn independently from the standard Gaussian distribution $\text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$, where the covariance $\mathbf{I}_d$ is the d -dimensional identity matrix. For this special case, we can obtain a sharp estimate for the minimum conic singular value $\lambda_{\min}(\Phi; K)$ for any convex cone K .

2.3.2 The conic Gaussian width

The analysis of Gaussian sampling depends on a geometric summary parameter for cones.

Definition 3.1 (Conic Gaussian width). Let $K \subset \mathbb{R}^d$ be a cone, not necessarily convex. The *conic Gaussian width* $w(K)$ is defined as

$$w(K) := \mathbb{E} \sup_{\mathbf{u} \in K \cap \mathbb{S}^{d-1}} \langle \mathbf{g}, \mathbf{u} \rangle$$

where $\mathbf{g} \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$ is a standard Gaussian vector in $\mathbb{R}^d$.

The Gaussian width plays a central role in asymptotic convex geometry [17, 23, 27]. Most of the classical techniques for bounding widths are only accurate up to constant factors (or worse). In contrast, ideas from the contemporary signal processing literature frequently allow us to produce numerically sharp estimates for the Gaussian width of a cone. These techniques were developed in the papers [1, 6, 10, 24, 31]. We will outline one of the methods in Section 2.4.

Remark 3.2 (Statistical dimension). The conic Gaussian width $w(K)$ is a convenient functional because it arises from the probabilistic tools that we use. The theory of conic integral geometry, however, delivers a better summary parameter [1]. The *statistical dimension* $\delta(K)$ of a convex cone K can be defined as

$$\delta(K) := \mathbb{E} \left[\left(\sup_{\mathbf{u} \in K \cap \mathbf{B}^d} \langle \mathbf{g}, \mathbf{u} \rangle \right)^2 \right],$$

where $\mathbf{B}^d$ is the Euclidean unit ball in $\mathbb{R}^d$ and $\mathbf{g} \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$. The statistical dimension canonically extends the dimension of a subspace to the class of convex cones, and it satisfies many elegant identities [1, Prop. 3.1]. For some purposes, the two parameters are interchangeable because of the following comparison [1, Prop. 10.2]:

$$w^2(K) \leq \delta(K) \leq w^2(K) + 1.$$

As a consequence, we can interpret $w^2(K)$ as a rough measure of the “dimension” of a cone.

2.3.3 Conic singular values and conic Gaussian widths

As it turns out, the conic Gaussian width $w(K)$ controls the minimum conic singular value $\lambda_{\min}(\Phi; K)$ when Φ follows the standard normal distribution.

Proposition 3.3 (Minimum conic singular value of a Gaussian matrix). *Let $K \subset \mathbb{R}^d$ be a cone, not necessarily convex, and let Φ be an $m \times d$ matrix whose rows are independent vectors drawn from the standard Gaussian distribution $\text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$. Then*

$$\lambda_{\min}(\Phi; K) \geq \sqrt{m-1} - w(K) - t$$

with probability at least $1 - e^{-t^2/2}$.

In essence, this result dates to the work of Gordon [11, 12]. We have drawn the proof from the survey [8, Sec. 3.2] of Davidson & Szarek; see also [6, 22, 29, 31]. Note that the argument relies on special results for Gaussian processes that do not extend to other distributions.

Proof sketch. We can express the minimum conic singular value as

$$\lambda_{\min}(\Phi; K) = \inf_{u \in K \cap S^{d-1}} \sup_{v \in S^{m-1}} \langle v, \Phi u \rangle$$

It is a consequence of Gordon's comparison inequality [11, Thm. 1.4] that

$$\mathbb{E} \inf_{u \in K \cap S^{d-1}} \sup_{v \in S^{m-1}} \langle v, \Phi u \rangle \geq \mathbb{E} \sup_{v \in S^{m-1}} \langle g', v \rangle - \mathbb{E} \sup_{u \in K \cap S^{d-1}} \langle g, u \rangle = \mathbb{E} \|g'\| - w(K),$$

where $g' \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_m)$ and $g \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$. It is well known that $\mathbb{E} \|g'\| \geq \sqrt{m-1}$, and therefore

$$\mathbb{E} \lambda_{\min}(\Phi; K) \geq \sqrt{m-1} - w(K). \quad (2.7)$$

To complete the argument, note that the map

$$\lambda_{\min}(\cdot; K) : A \mapsto \inf_{u \in K \cap S^{d-1}} \|Au\|$$

is 1-Lipschitz with respect to the Frobenius norm. The usual Gaussian concentration inequality [3, Sec. 5.4] implies that

$$\mathbb{P}\{\lambda_{\min}(\Phi; K) \leq \mathbb{E} \lambda_{\min}(\Phi; K) - t\} \leq e^{-t^2/2}. \quad (2.8)$$

Introduce the lower bound (2.7) for the expectation of the minimum conic singular value into (2.8) to reach the advertised result. $\square$

Remark 3.4 (Sharpness for convex cones). It is a remarkable fact that the bound in Proposition 3.3 is essentially sharp. For any cone K , we can reinterpret the statement as saying that

$$\lambda_{\min}(\Phi; K) > 0 \quad \text{with high probability when} \quad m \geq w^2(K) + Cw(K).$$

(The letter C always denotes a positive absolute constant, but its value may change from place to place.) Conversely, for a convex cone K , it holds that

$$\lambda_{\min}(\Phi; K) = 0 \quad \text{with high probability when} \quad m \leq w^2(K) - Cw(K). \quad (2.9)$$

Result (2.9) follows from research of Amelunxen et al. [1, Thm. I and Prop. 10.2]. This claim can also be derived by supplementing the proof of Proposition 3.3 with a short polarity argument. It is productive to interpret the pair of estimates in this remark as a *phase transition* for convex signal recovery; see [1] for more information.

2.3.4 An error bound for Gaussian measurements

Combining Proposition 2.6 and Proposition 3.3, we obtain a general error bound for convex recovery from Gaussian measurements.

Corollary 3.5 (Signal recovery from Gaussian measurements). *Let $\mathbf{x}^\natural$ be a signal in $\mathbb{R}^d$. Let Φ be an $m \times d$ matrix whose rows are independent random vectors drawn from the standard Gaussian distribution $\text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$, and let $\mathbf{y} = \Phi \mathbf{x}^\natural + \mathbf{e}$ be a vector of measurements in $\mathbb{R}^m$. With probability at least $1 - e^{-t^2/2}$, the following statement holds. Assume that $\|\mathbf{e}\| \leq \eta$, and let $\hat{\mathbf{x}}_\eta$ be any solution to the optimization problem (2.2). Then*

$$\|\hat{\mathbf{x}}_\eta - \mathbf{x}^\natural\| \leq \frac{2\eta}{[\sqrt{m-1} - w(\mathcal{D}(f, \mathbf{x}^\natural)) - t]_+}.$$

The operation $[a]_+ := \max\{a, 0\}$ returns the positive part of a number.

The overall argument that leads to this result was proposed by Rudelson & Vershynin [29, Sec. 4]; the statement here is adapted from [6].

Corollary 3.5 provides for stable recovery of the unknown $\mathbf{x}^\natural$ when the number m of measurements satisfies

$$m \geq w^2(\mathcal{D}(f, \mathbf{x}^\natural)) + Cw(\mathcal{D}(f, \mathbf{x}^\natural)).$$

In view of Remark 3.4, Corollary 3.5 provides a refined estimate for the amount of information that suffices to identify a structured vector from Gaussian measurements via convex optimization.

Remark 3.6 (The normal error model). It is possible to improve the error bound in Corollary 3.5 if we instate a Gaussian model for the error vector $\mathbf{e}$. See the papers [25, 26, 33] for an analysis of this case.

2.4 Controlling the width of a descent cone via polarity

As soon as we know the conic Gaussian width of the descent cone, Corollary 6.4 yields error bounds for convex recovery of a structured signal from Gaussian measurements. To make use of this result, we need technology for calculating these widths. This section describes a mechanism, based on polarity, that leads to extremely accurate estimates. We can trace this method to the papers [24, 31], where it is couched in the language of duality for cone programs. The subsequent papers [1, 6] rephrase these ideas in a more geometric fashion. It can be shown that the approach in this section gives sharp results for many natural examples; see [1, Thm. 4.3] or [10, Prop. 1]. Although polar bounds for widths are classic in asymptotic convex geometry [17, 23, 27], the refined arguments here are just a few years old.

2.4.1 Polarity and weak duality for cones

We begin with some classical facts about conic geometry.

Fact 4.1 (Polarity). *Let K be a general cone in $\mathbb{R}^d$. The polar cone K° is the closed convex cone*

$$K^\circ := \{\mathbf{v} \in \mathbb{R}^d : \langle \mathbf{v}, \mathbf{x} \rangle \leq 0 \text{ for all } \mathbf{x} \in K\}.$$

It is easy to verify that $K \subset (K^\circ)^\circ$.

Recall that the *distance* from a point $\mathbf{x} \in \mathbb{R}^d$ to a set $E \subset \mathbb{R}^d$ is defined by the relation

$$\text{dist}(\mathbf{x}, E) := \inf_{\mathbf{u} \in E} \|\mathbf{x} - \mathbf{u}\|.$$

With these definitions, we reach the following weak duality result.

Proposition 4.2 (Weak duality for cones). *Let K be a general cone in $\mathbb{R}^d$. For $\mathbf{x} \in \mathbb{R}^d$,*

$$\sup_{\mathbf{u} \in K \cap \mathbb{S}^{d-1}} \langle \mathbf{x}, \mathbf{u} \rangle \leq \text{dist}(\mathbf{x}, K^\circ).$$

Proof. The argument is based on a simple duality trick. First, write

$$\text{dist}(\mathbf{x}, K^\circ) = \inf_{\mathbf{v} \in K^\circ} \|\mathbf{x} - \mathbf{v}\| = \inf_{\mathbf{v} \in K^\circ} \sup_{\mathbf{u} \in \mathbb{S}^{d-1}} \langle \mathbf{x} - \mathbf{v}, \mathbf{u} \rangle.$$

Apply the inf–sup inequality:

$$\text{dist}(\mathbf{x}, K^\circ) \geq \sup_{\mathbf{u} \in \mathbb{S}^{d-1}} \inf_{\mathbf{v} \in K^\circ} \langle \mathbf{x} - \mathbf{v}, \mathbf{u} \rangle = \sup_{\mathbf{u} \in \mathbb{S}^{d-1}} \left[\langle \mathbf{x}, \mathbf{u} \rangle - \sup_{\mathbf{v} \in K^\circ} \langle \mathbf{v}, \mathbf{u} \rangle \right].$$

By definition of polarity, the inner supremum takes the value $+\infty$ unless $\mathbf{u} \in (K^\circ)^\circ$. We determine that

$$\text{dist}(\mathbf{x}, K^\circ) \geq \sup_{\mathbf{u} \in (K^\circ)^\circ \cap \mathbb{S}^{d-1}} \langle \mathbf{x}, \mathbf{u} \rangle \geq \sup_{\mathbf{u} \in K \cap \mathbb{S}^{d-1}} \langle \mathbf{x}, \mathbf{u} \rangle.$$

The last inequality holds because $K \subset (K^\circ)^\circ$. $\square$

Remark 4.3 (Strong duality for cones). If K is a convex cone and we replace the sphere with a ball, then we have strong duality instead:

$$\sup_{\mathbf{u} \in K \cap \mathbb{B}^d} \langle \mathbf{x}, \mathbf{u} \rangle = \text{dist}(\mathbf{x}, K^\circ).$$

The proof uses Sion’s minimax theorem [30] and the bipolar theorem [28, Thm. 14.1].

2.4.2 The conic Gaussian width of a descent cone

We can use Proposition 4.2 to obtain an effective bound for the width of a descent cone. This approach is based on a classical polarity correspondence [28, Thm. 23.7].

Fact 4.4 (Polarity for descent cones). *The subdifferential of a proper convex function $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ at a point $\mathbf{x} \in \mathbb{R}^d$ is the closed convex set*

$$\partial f(\mathbf{x}) := \{\mathbf{v} \in \mathbb{R}^d : f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \mathbf{v}, \mathbf{y} - \mathbf{x} \rangle \text{ for all } \mathbf{y} \in \mathbb{R}^d\}.$$

Assume that the subdifferential $\partial f(\mathbf{x})$ is nonempty and does not contain the origin. Then

$$\mathcal{D}(f, \mathbf{x})^\circ = \overline{\text{cone}}(\partial f(\mathbf{x})) := \text{closure} \left(\bigcup_{\tau \geq 0} \tau \cdot \partial f(\mathbf{x}) \right). \quad (2.10)$$

Combining Proposition 4.2 and Fact 4.4, we reach a bound for the conic Gaussian width of a descent cone.

Proposition 4.5 (The width of a descent cone). *Let $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ be a proper convex function, and fix a point $\mathbf{x} \in \mathbb{R}^d$. Assume that the subdifferential $\partial f(\mathbf{x})$ is nonempty and does not contain the origin. Then*

$$w^2(\mathcal{D}(f, \mathbf{x})) \leq \mathbb{E} \inf_{\tau \geq 0} \text{dist}^2(\mathbf{g}, \tau \cdot \partial f(\mathbf{x}))$$

Several specific instances of Proposition 4.5 appear in [6, App. C], while the general statement here is adapted from [1, Sec. 4.1]. Sections 2.4.3 and 2.4.4 exhibit how Proposition 4.5 works.

Proof. Proposition 4.2 implies that

$$w(\mathcal{D}(f, \mathbf{x})) = \mathbb{E} \sup_{\mathbf{u} \in \mathcal{D}(f, \mathbf{x}) \cap \mathbb{S}^{d-1}} \langle \mathbf{g}, \mathbf{u} \rangle \leq \mathbb{E} \text{dist}(\mathbf{g}, \mathcal{D}(f, \mathbf{x})^\circ).$$

Expression (2.10) for the polar of a descent cone implies that

$$w(\mathcal{D}(f, \mathbf{x})) \leq \mathbb{E} \text{dist}\left(\mathbf{g}, \text{closure}\left(\bigcup_{\tau \geq 0} \tau \cdot \partial f(\mathbf{x})\right)\right) = \mathbb{E} \inf_{\tau \geq 0} \text{dist}(\mathbf{g}, \tau \cdot \partial f(\mathbf{x})).$$

Indeed, the distance to a set is the same as the distance to its closure, and the distance to a union is the infimal distance to one of its members. Square the latter display and apply Jensen's inequality to complete the argument. $\square$

2.4.3 Example: Sparse vectors

Suppose that $\mathbf{x}^\natural$ is a vector in $\mathbb{R}^d$ with s nonzero entries. Let Φ be an $m \times d$ matrix whose rows are independent random vectors distributed as $\text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$, and suppose that we acquire a vector $\mathbf{y} = \Phi \mathbf{x}^\natural + \mathbf{e}$ consisting of m noisy measurements. We can solve the ℓ_1 -minimization problem (2.3) in an attempt to reconstruct $\mathbf{x}^\natural$.

How many measurements are sufficient to ensure that this approach succeeds? We will demonstrate that

$$w^2(\mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{x}^\natural)) \leq 2s \log(d/s) + 2s. \quad (2.11)$$

Therefore, Corollary 3.5 implies that $m \gtrsim 2s \log(d/s)$ measurements are enough for us to recover $\mathbf{x}^\natural$ approximately. When $s \ll d$, the first term in (2.11) is numerically sharp because of [10, Prop. 1].

2.4.3.1 The width calculation

Let us establish the width bound (2.11). This analysis is adapted from [6, App. C] and [1, App. D.2]; see also [10, App. B]. The result [1, Prop. 4.5] contains a more complicated formula for the width that is sharp for all choices of the sparsity s .

When estimating widths, a useful strategy is to change coordinates so that the calculations are more transparent. The ℓ_1 norm is invariant under signed permutation, so

$$\mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{x}^\natural) = \mathbf{P}^\mathbf{t} \mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{P}\mathbf{x}^\natural) \quad \text{where } \mathbf{P} \text{ is a signed permutation.}$$

The distribution of a standard Gaussian random variable is invariant under signed permutation, so the conic Gaussian width has the same invariance. Therefore,

$$w(\mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{x}^\natural)) = w(\mathbf{P}^\mathbf{t} \mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{P}\mathbf{x}^\natural)) = w(\mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{P}\mathbf{x}^\natural)).$$

We will use this type of transformation several times without detailed justification.

As a consequence of the argument in the last paragraph, we may assume that $\mathbf{x}^\natural$ takes the form

$$\mathbf{x}^\natural = (x_1, \dots, x_s, 0, \dots, 0)^\mathbf{t} \in \mathbb{R}^d \quad \text{where } x_1 \geq \dots \geq x_s > 0.$$

Proposition 4.5 ensures that

$$w^2(\mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{x}^\natural)) \leq \mathbb{E} \operatorname{dist}^2(\mathbf{g}, \tau \cdot \partial \|\mathbf{x}^\natural\|_{\ell_1}) \quad \text{for each } \tau \geq 0 \quad (2.12)$$

where $\mathbf{g} \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$. The subdifferential of the ℓ_1 norm at $\mathbf{x}^\natural$ satisfies

$$\partial \|\mathbf{x}^\natural\|_{\ell_1} = \left\{ \begin{bmatrix} \mathbf{1}_s \\ \mathbf{y} \end{bmatrix} \in \mathbb{R}^d : \|\mathbf{y}\|_{\ell_\infty} \leq 1 \right\} \quad \text{where } \mathbf{1}_s := (1, \dots, 1)^\mathbf{t} \in \mathbb{R}^s.$$

Therefore,

$$\mathbb{E} \operatorname{dist}^2(\mathbf{g}, \tau \cdot \partial \|\mathbf{x}^\natural\|_{\ell_1}) = \sum_{j=1}^s \mathbb{E} (g_j - \tau)^2 + \sum_{j=s+1}^d \mathbb{E} [|g_j| - \tau]_+^2. \quad (2.13)$$

As usual, $[a]_+ := \max\{a, 0\}$. For $1 \leq j \leq s$, a direct calculation gives

$$\mathbb{E} (g_j - \tau)^2 = 1 + \tau^2. \quad (2.14)$$

For $s < j \leq d$, we apply a familiar tail bound for the standard normal variable to obtain

$$\begin{aligned} \mathbb{E} [|g_j|^2 - \tau]_+^2 &= \int_\tau^\infty 2(a - \tau)^2 \mathbb{P}\{|g_j| \geq a\} da \\ &\leq \int_\tau^\infty 2a \left(\sqrt{\frac{2}{\pi}} a^{-1} e^{-a^2/2} \right) da \\ &= 2\mathbb{P}\{|g_j| \geq \tau\} \leq 2e^{-\tau^2/2} \end{aligned} \quad (2.15)$$

Combine (2.12), (2.13), (2.14), and (2.15) to obtain

$$w^2(\mathcal{D}(\|\cdot\|_{\ell_1}, \mathbf{x}^\natural)) \leq \mathbb{E} \operatorname{dist}^2(\mathbf{g}, \tau \cdot \partial \|\mathbf{x}^\natural\|_{\ell_1}) = s \cdot (1 + \tau^2) + (d - s) \cdot 2e^{-\tau^2/2}.$$

Choose $\tau^2 = 2 \log(d/s)$ and simplify to reach (2.11).

2.4.4 Example: Low-rank matrices

Let $\mathbf{X}^\natural$ be a matrix in $\mathbb{R}^{d_1 \times d_2}$ with rank r . Let $\Phi : \mathbb{R}^{d_1 \times d_2} \rightarrow \mathbb{R}^m$ be a linear operator whose matrix has independent standard Gaussian entries. Suppose we acquire m noisy measurements of the form $\mathbf{y} = \Phi(\mathbf{X}^\natural) + \mathbf{e}$. We can solve the S_1 -minimization problem (2.5) to reconstruct $\mathbf{X}^\natural$.

How many measurements are enough to guarantee that this approach works? We will prove that

$$w^2(\mathcal{D}(\|\cdot\|_{S_1}, \mathbf{X}^\natural)) \leq 3r \cdot (d_1 + d_2 - r). \quad (2.16)$$

As a consequence, Corollary 3.5 implies that $m \gtrsim 3r \cdot (d_1 + d_2 - r)$ measurements allow us to identify $\mathbf{X}^\natural$ approximately.

2.4.4.1 The width calculation

Let us establish the width bound (2.16). This analysis is adapted from [6, App. C] and [1, App. D.3]; see also [10, App. E]. The result [1, Prop. 4.6] contains a more complicated formula for the width that is sharp whenever the rank r is proportional to the dimension $\min\{d_1, d_2\}$.

The Schatten 1-norm is unitarily invariant, so we may also select a coordinate system where

$$\mathbf{X}^\natural = \begin{bmatrix} \Sigma & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \quad \text{where} \quad \Sigma = \operatorname{diag}(\sigma_1, \dots, \sigma_r) \quad \text{and} \quad \sigma_j > 0 \text{ for } j = 1, \dots, r.$$

Let $\mathbf{G}$ be a $d_1 \times d_2$ matrix with independent standard normal entries, partitioned as

$$\mathbf{G} = \begin{bmatrix} \mathbf{G}_{11} & \mathbf{G}_{12} \\ \mathbf{G}_{21} & \mathbf{G}_{22} \end{bmatrix} \quad \text{where} \quad \mathbf{G}_{11} \text{ is } r \times r \quad \text{and} \quad \mathbf{G}_{22} \text{ is } (d_1 - r) \times (d_2 - r).$$

Define a random parameter $\tau = \|\mathbf{G}_{22}\|$, where $\|\cdot\|$ denotes the spectral norm. Proposition 4.5 ensures that

$$w^2(\mathcal{D}(\|\cdot\|_{S_1}, \mathbf{X}^\natural)) \leq \mathbb{E} \operatorname{dist}_{\mathbb{F}}^2(\mathbf{G}, \tau \cdot \partial \|\mathbf{X}^\natural\|_{S_1}). \quad (2.17)$$

Note that we must calculate distance with respect to the Frobenius norm $\|\cdot\|_F$. According to [37, Ex. 2], the subdifferential of the Schatten 1-norm takes the form

$$\partial \|X^\natural\|_{S_1} = \left\{ \begin{bmatrix} \mathbf{I}_r & \mathbf{0} \\ \mathbf{0} & \mathbf{Y} \end{bmatrix} \in \mathbb{R}^{d_1 \times d_2} : \|\mathbf{Y}\| \leq 1 \right\} \quad \text{where } \mathbf{I}_r \text{ is the } r \times r \text{ identity matrix.}$$

We may calculate that

$$\begin{aligned} \mathbb{E} \text{dist}_F^2(\mathbf{G}, \tau \cdot \|X^\natural\|_{S_1}) &= \mathbb{E} \|\mathbf{G}_{11} - \tau \cdot \mathbf{I}_r\|_F^2 + \mathbb{E} \|\mathbf{G}_{12}\|_F^2 \\ &\quad + \mathbb{E} \|\mathbf{G}_{21}\|_F^2 + \mathbb{E} \inf_{\|\mathbf{Y}\| \leq 1} \|\mathbf{G}_{22} - \tau \cdot \mathbf{Y}\|_F^2. \end{aligned} \quad (2.18)$$

Our selection of τ ensures that the last term on the right-hand side of (2.18) vanishes. By direct calculation,

$$\mathbb{E} \|\mathbf{G}_{12}\|_F^2 + \mathbb{E} \|\mathbf{G}_{21}\|_F^2 = r \cdot (d_1 + d_2 - 2r). \quad (2.19)$$

To bound the first term on right-hand side of (2.18), observe that

$$\mathbb{E} \|\mathbf{G}_{11} - \tau \cdot \mathbf{I}_r\|_F^2 = r^2 + r \cdot \mathbb{E} \tau^2 \quad (2.20)$$

because the random variable τ is independent of $\mathbf{G}_{11}$. We need to compute $\mathbb{E} \tau^2 = \mathbb{E} \|\mathbf{G}_{22}\|_F^2$. A short argument [8, Sec. 2.3] based on the Slepian comparison inequality shows that

$$\mathbb{E} \|\mathbf{G}_{22}\| \leq \sqrt{d_1 - r} + \sqrt{d_2 - r} \leq \sqrt{2(d_1 + d_2 - 2r)}.$$

The spectral norm is 1-Lipschitz, so the Gaussian Poincaré inequality [3, Thm. 3.20] implies

$$\mathbb{E} \|\mathbf{G}_{22}\|^2 - (\mathbb{E} \|\mathbf{G}_{22}\|)^2 = \text{Var}(\|\mathbf{G}_{22}\|) \leq 1.$$

Combining the last two displays,

$$\mathbb{E} \tau^2 = \mathbb{E} \|\mathbf{G}_{22}\|^2 \leq (\mathbb{E} \|\mathbf{G}_{22}\|)^2 + 1 \leq 2(d_1 + d_2 - 2r) + 1. \quad (2.21)$$

Finally, we incorporate (2.18), (2.19), (2.20), (2.21) into the width bound (2.17) to reach

$$w^2(\mathcal{D}(\|\cdot\|_{S_1}, X^\natural)) \leq 3r \cdot (d_1 + d_2 - 2r) + r^2 + r.$$

Simplify this expression to obtain result (2.16).

2.5 Mendelson's Small Ball Method

In Sections 2.2–2.4, we analyzed a convex programming method for recovering structured signals from standard Gaussian measurements. The main result, Corollary 3.5, is appealing because it applies to any convex complexity measure f . Proposition 4.5 allows us to instantiate this result because it provides a mechanism for controlling the Gaussian width of a descent cone. On the other hand, this approach only works when the sampling matrix Φ follows the standard Gaussian distribution.

For other sampling models, researchers use a variety of ad hoc techniques to study the recovery problem. It is common to see a separate and intricate argument for each new complexity measure f and each new distribution for Φ . It is natural to wonder whether there is a single approach that can address a broad class of complexity measures and sampling matrices.

The primary goal of this chapter is to analyze convex signal reconstruction with more general random measurements. Our argument is based on Mendelson's *Small Ball Method*, a powerful strategy for establishing a lower bound on a nonnegative empirical process [14, 16, 19–21]. This section contains an overview of Mendelson's Small Ball Method. Section 2.6 uses this technique to study sub-Gaussian measurement models. In Section 2.7, we extend these ideas to a larger class of sampling distributions. In Section 2.8, we conclude with an application to the problem of phase retrieval.

2.5.1 The minimum conic singular value as a nonnegative empirical process

Suppose that φ is a random vector on $\mathbb{R}^d$, and draw independent copies $\varphi_1, \dots, \varphi_m$ of the random vector φ . Form an $m \times d$ sampling matrix Φ whose rows are these random vectors:

$$\Phi = \begin{bmatrix} \varphi_1^\top \\ \vdots \\ \varphi_m^\top \end{bmatrix}. \quad (2.22)$$

Fix a cone $K \in \mathbb{R}^d$, not necessarily convex, and define the set $E := K \cap \mathbb{S}^{d-1}$. Then we can express the minimum conic singular value $\lambda_{\min}(\Phi; K)$ of the sampling matrix as a nonnegative empirical process:

$$\lambda_{\min}(\Phi; K) = \inf_{u \in E} \left(\sum_{i=1}^m |\langle \varphi_i, u \rangle|^2 \right)^{1/2}. \quad (2.23)$$

When the sampling matrix is Gaussian, we can use Gordon's theorem [11, Thm. 1.4] to obtain a lower bound for expression (2.23), as in Proposition 3.3. The challenge is to find an alternative method for producing a lower bound in a more general setting.

2.5.2 A lower bound for nonnegative empirical processes

The main technical component in Mendelson's Small Ball Method is a remarkable estimate that was developed in the paper [20]. This result delivers an effective lower bound for a nonnegative empirical process.

Proposition 5.1 (Lower bound for a nonnegative empirical process [20, Thm. 5.4]). *Fix a set $E \subset \mathbb{R}^d$. Let $\boldsymbol{\varphi}$ be a random vector on $\mathbb{R}^d$, and let $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_m$ be independent copies of $\boldsymbol{\varphi}$. Define the $m \times d$ matrix Φ as on (2.22). Introduce the marginal tail function*

$$Q_\xi(E; \boldsymbol{\varphi}) := \inf_{\mathbf{u} \in E} \mathbb{P}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq \xi\} \quad \text{where } \xi \geq 0.$$

Let $\varepsilon_1, \dots, \varepsilon_m$ be independent Rademacher random variables,² independent of everything else, and define the mean empirical width of the set:

$$W_m(E; \boldsymbol{\varphi}) := \mathbb{E} \sup_{\mathbf{u} \in E} \langle \mathbf{h}, \mathbf{u} \rangle \quad \text{where} \quad \mathbf{h} := \frac{1}{\sqrt{m}} \sum_{i=1}^m \varepsilon_i \boldsymbol{\varphi}_i. \quad (2.24)$$

Then, for any $\xi > 0$ and $t > 0$,

$$\inf_{\mathbf{u} \in E} \left(\sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle|^2 \right)^{1/2} \geq \xi \sqrt{m} Q_{2\xi}(E; \boldsymbol{\varphi}) - 2W_m(E; \boldsymbol{\varphi}) - \xi t$$

with probability at least $1 - e^{-t^2/2}$.

The proof appears below in Section 2.5.5. In the sequel, we usually lighten our notation for Q_ξ and W_m by suppressing the dependence on $\boldsymbol{\varphi}$.

Before we continue, it may be helpful to remark on this result. The marginal tail function $Q_\xi(E)$ reflects the probability that the random variable $|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle|$ is close to zero for any fixed vector $\mathbf{u} \in E$. When $Q_\xi(E)$ is bounded away from zero for some ξ , the nonnegative empirical process is likely to be large. Koltchinskii & Mendelson [14] point out that the marginal tail function reflects the absolute continuity of the distribution of $\boldsymbol{\varphi}$, so Q_ξ may be quite small when $\boldsymbol{\varphi}$ is “spiky.”

²A Rademacher random variable takes the two values ± 1 with equal probability.

The mean empirical width $W_m(E)$ is a distribution-dependent measure of the size of the set E . When $\boldsymbol{\varphi}$ follows a standard Gaussian distribution, $W_m(E)$ reduces to the usual Gaussian width $W(E) := \mathbb{E} \sup_{\mathbf{u} \in E} \langle \mathbf{g}, \mathbf{u} \rangle$. As the number m tends to infinity, the distribution of the random vector $\mathbf{h}$ converges in distribution to a centered Gaussian variable with covariance $\mathbb{E}[\boldsymbol{\varphi}\boldsymbol{\varphi}^*]$. Therefore, $W_m(E) \rightarrow W(E)$ when $\boldsymbol{\varphi}$ is centered and isotropic.

2.5.3 Mendelson's Small Ball Method

Proposition 5.1 shows that we can obtain a lower bound for (2.23) by performing two simpler estimates. To achieve this goal, Mendelson has developed a general strategy, which consists of three steps:

MENDELSON'S SMALL BALL METHOD

- (1) Apply Proposition 5.1 to bound the minimum conic singular value $\lambda_{\min}(\boldsymbol{\Phi}; K)$ below in terms of the marginal tail function $Q_{2\xi}(E; \boldsymbol{\varphi})$ and the mean empirical width $W_m(E; \boldsymbol{\varphi})$. The index set $E := K \cap \mathbb{S}^{d-1}$.
- (2) Bound the marginal tail function $Q_{2\xi}(E; \boldsymbol{\varphi})$ below using a Paley–Zygmund inequality.
- (3) Bound the mean empirical width $W_m(E; \boldsymbol{\varphi})$ above by imitating techniques for controlling the Gaussian width of E .

This presentation is distilled from the corpus [14, 16, 19–21]. A more sophisticated variant of this method appears in [20, Thm. 5.3]. Later in this chapter, we will encounter several concrete applications of this strategy.

2.5.4 Expected Scope

Mendelson's Small Ball Method provides lower bounds for (2.23) in many situations, but it does not offer a universal prescription. Let us try to delineate the circumstances where this approach is likely to be useful for signal recovery problems.

- Mendelson's Small Ball Method assumes that the sampling matrix $\boldsymbol{\Phi}$ has independent, identically distributed rows. Although this model describes many of the sampling strategies in the literature, there are some examples, such as random filtering [34], that do not conform to this assumption.

- A major advantage of Mendelson's Small Ball Method is that it applies to sampling distributions with heavy tails. On the other hand, the random vector $\boldsymbol{\varphi}$ cannot be too "spiky," or else it may not be possible to produce a good lower bound for the marginal tail function $Q_{2\xi}(E)$. This requirement indicates that the approach may require significant improvements before it applies to problems like matrix completion.

There are a number of possible extensions of Mendelson's Small Ball Method that could expand its bailiwick. For example, it is easy to extend Proposition 5.1 to address the case where the random vector $\boldsymbol{\varphi}$ is complex valued. A more difficult, but very useful, modification would allow us to block the measurements into groups. This revision could reduce the difficulties associated with spiky distributions, but it seems to demand some additional ideas.

2.5.5 Proof of Proposition 5.1

Let us establish the Mendelson bound for a nonnegative empirical process. First, we introduce a directional version of the marginal tail function:

$$Q_{\xi}(\mathbf{u}) := \mathbb{P}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq \xi\} \quad \text{for } \mathbf{u} \in E \text{ and } \xi > 0.$$

Lyapunov's inequality and Markov's inequality give the numerical bounds

$$\left(\frac{1}{m} \sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle|^2 \right)^{1/2} \geq \frac{1}{m} \sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle| \geq \frac{\xi}{m} \sum_{i=1}^m \mathbb{1}\{|\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle| \geq \xi\}.$$

We write $\mathbb{1}A$ for the 0–1 random variable that takes the value one when the event A takes place. Add and subtract $Q_{2\xi}(\mathbf{u})$ inside the sum and then take the infimum over $\mathbf{u} \in E$ to reach the inequality

$$\begin{aligned} \inf_{\mathbf{u} \in E} \left(\frac{1}{m} \sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle|^2 \right)^{1/2} &\geq \xi \inf_{\mathbf{u} \in E} Q_{2\xi}(\mathbf{u}) \\ &\quad - \frac{\xi}{m} \sup_{\mathbf{u} \in E} \sum_{i=1}^m [Q_{2\xi}(\mathbf{u}) - \mathbb{1}\{|\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle| \geq \xi\}]. \end{aligned} \tag{2.25}$$

To control the supremum in probability, we can invoke the bounded difference inequality [3, Sec. 6.1]. Observe that each summand is independent and bounded in magnitude by one. Therefore,

$$\begin{aligned}
& \sup_{\mathbf{u} \in E} \sum_{i=1}^m [Q_{2\xi}(\mathbf{u}) - \mathbb{1}\{|\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle| \geq \xi\}] \\
& \leq \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m [Q_{2\xi}(\mathbf{u}) - \mathbb{1}\{|\langle \boldsymbol{\varphi}_i, \mathbf{x} \rangle| \geq \xi\}] + t\sqrt{m} \quad (2.26)
\end{aligned}$$

with probability at least $1 - e^{-t^2/2}$.

Next, we simplify the expected supremum. Introduce a soft indicator function:

$$\psi_\xi : \mathbb{R} \rightarrow [0, 1] \quad \text{where} \quad \psi_\xi(s) := \begin{cases} 0, & |s| \leq \xi \\ (|s| - \xi)/\xi, & \xi < |s| \leq 2\xi \\ 1, & 2\xi < |s|. \end{cases}$$

We need two properties of the soft indicator. First, the soft indicator is bracketed by two hard indicators: $\mathbb{1}\{|s| \geq 2\xi\} \leq \psi_\xi(s) \leq \mathbb{1}\{|s| \geq \xi\}$ for all $s \in \mathbb{R}$. Second, $\xi\psi_\xi$ is a *contraction*, i.e., a 1-Lipschitz function on $\mathbb{R}$ that fixes the origin. Therefore, we can make the following calculation:

$$\begin{aligned}
& \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m [Q_{2\xi}(\mathbf{u}) - \mathbb{1}\{|\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle| \geq \xi\}] \\
& = \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m [\mathbb{E} \mathbb{1}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq 2\xi\} - \mathbb{1}\{|\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle| \geq \xi\}] \\
& \leq \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m [\mathbb{E} \psi_\xi(\langle \boldsymbol{\varphi}, \mathbf{u} \rangle) - \psi_\xi(\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle)] \\
& \leq 2 \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m \varepsilon_i \psi_\xi(\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle) \\
& \leq \frac{2}{\xi} \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m \varepsilon_i \langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle. \quad (2.27)
\end{aligned}$$

In the first line, we write the marginal tail function as an expectation and then we bound the two indicators using the soft indicator function. The next inequality is the Giné–Zinn symmetrization [35, Lem. 2.3.1]. The last line follows from the Rademacher comparison principle [17, Eqn. (4.20)] because $\xi\psi_\xi$ is a contraction.

Combine the inequalities (2.25), (2.26), and (2.27) to reach

$$\inf_{\mathbf{u} \in E} \left(\frac{1}{m} \sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle|^2 \right)^{1/2} \geq \xi \inf_{\mathbf{u} \in E} Q_{2\xi}(\mathbf{u}) - \frac{\xi}{m} \left[\frac{2}{\xi} \mathbb{E} \sup_{\mathbf{u} \in E} \sum_{i=1}^m \varepsilon_i \langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle + t\sqrt{m} \right].$$

Define $\mathbf{h} := m^{-1/2} \sum_{i=1}^m \varepsilon_i \boldsymbol{\varphi}_i$ and clear the factor $\sqrt{m}$ to conclude that

$$\inf_{\mathbf{u} \in E} \left(\sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle|^2 \right)^{1/2} \geq \xi \sqrt{m} \inf_{\mathbf{u} \in E} Q_{2\xi}(\mathbf{u}) - 2 \mathbb{E} \sup_{\mathbf{u} \in E} \langle \mathbf{h}, \mathbf{u} \rangle - \xi t.$$

with probability at least $1 - e^{-t^2/2}$. Identify the marginal tail function $Q_{2\xi}(E)$ and the empirical width $W_m(E)$ to establish Proposition 5.1.

2.6 A universal error bound for sub-Gaussian measurements

In this section, we invoke Mendelson's Small Ball Method to study convex signal recovery from independent sub-Gaussian measurements. This class of examples provides a wide generalization of standard Gaussian measurements. We will establish a variant of the Gaussian recovery result, Corollary 3.5, in this setting.

2.6.1 Sub-Gaussian measurements

Let us set out the conditions we require for the sampling matrix. Suppose that $\boldsymbol{\varphi}$ is a random vector in $\mathbb{R}^d$ that has the following properties:

- **[Centering]** The vector has zero mean: $\mathbb{E} \boldsymbol{\varphi} = \mathbf{0}$.
- **[Nondegeneracy]** There is a positive constant α for which

$$\alpha \leq \mathbb{E} |\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \quad \text{for each } \mathbf{u} \in \mathbb{S}^{d-1}.$$

- **[Sub-Gaussian marginals]** There is a positive constant σ for which

$$\mathbb{P}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq t\} \leq 2e^{-t^2/(2\sigma^2)} \quad \text{for each } \mathbf{u} \in \mathbb{S}^{d-1}.$$

- **[Low eccentricity]** The eccentricity $\rho := \sigma/\alpha$ of the distribution should be small.

Finally, we construct a random $m \times d$ sampling matrix Φ whose rows are independent copies of $\boldsymbol{\varphi}^\dagger$, as in expression (2.22).

A few examples of sub-Gaussian distributions may be helpful.

Example 6.1 (Nonstandard Gaussian matrices). Suppose that $\boldsymbol{\varphi} \in \mathbb{R}^d$ follows the $\text{NORMAL}(\mathbf{0}, \Sigma)$ distribution where the covariance Σ satisfies $\frac{\pi}{2}\alpha^2 \leq \mathbf{u}^\dagger \Sigma \mathbf{u} \leq \sigma^2$ for each vector $\mathbf{u} \in \mathbb{S}^{d-1}$. Then the required conditions follow from basic facts about a normal distribution.

Example 6.2 (Independent bounded entries). Let X be a symmetric random variable whose magnitude is bounded by σ . Suppose that each entry of $\boldsymbol{\varphi}$ is an independent copy of X .

The vector $\boldsymbol{\varphi}$ inherits centering from X . Next, $\boldsymbol{\varphi}$ is nondegenerate with $\alpha \geq 2^{-1/2} \mathbb{E}|X|$ because of the Khintchine inequality [15] and a convexity argument. Finally, $\boldsymbol{\varphi}$ has sub-Gaussian marginals with the parameter σ because of Hoeffding's inequality [3, Sec. 2.6].

2.6.2 The minimum conic singular value of a sub-Gaussian matrix

The main result of this section gives a lower bound for the minimum conic singular value of a matrix Φ that satisfies the conditions in Section 2.6.1.

Theorem 6.3 (Minimum conic singular value of a sub-Gaussian matrix). *Suppose Φ is an $m \times d$ random matrix that satisfies the conditions in Section 2.6.1. Let $K \subset \mathbb{R}^d$ be a cone, not necessarily convex. Then*

$$\lambda_{\min}(\Phi; K) \geq c\alpha\rho^{-2} \cdot \sqrt{m} - C\sigma \cdot w(K) - \alpha t$$

with probability at least $1 - e^{-ct^2}$. The quantities c and C are positive absolute constants.

Observe that, when the eccentricity ρ has constant order, the bound in Theorem 6.3 matches the result for Gaussian matrices in Proposition 3.3. A similar result appears in the paper [22], so we do not claim any novelty. We establish Theorem 6.3 below in Section 2.6.4.

2.6.3 An error bound for sub-Gaussian measurements

Combining Proposition 2.6 and Theorem 6.3, we reach an immediate consequence for signal recovery from sub-Gaussian measurements.

Corollary 6.4 (Signal recovery from sub-Gaussian measurements). *Let $\mathbf{x}^\natural$ be a signal in $\mathbb{R}^d$. Let Φ be an $m \times d$ random matrix that satisfies the conditions in Section 2.6.1, and let $\mathbf{y} = \Phi\mathbf{x}^\natural + \mathbf{e}$ be a vector of measurements in $\mathbb{R}^m$. With probability at least $1 - e^{-ct^2}$, the following statement holds. Assume that $\|\mathbf{e}\| \leq \eta$, and let $\hat{\mathbf{x}}_\eta$ be any solution to the optimization problem (2.2). Then*

$$\|\hat{\mathbf{x}}_\eta - \mathbf{x}^\natural\| \leq \frac{2\eta}{[c\alpha\rho^{-2} \cdot \sqrt{m} - C\sigma \cdot w(\mathcal{D}(f, \mathbf{x}^\natural)) - \alpha t]_+}.$$

The quantities c and C are positive absolute constants. The operation $[a]_+ := \max\{a, 0\}$ returns the positive part of a number.

Corollary 6.4 provides for stable recovery of $\mathbf{x}^\natural$ as soon as the number m of sub-Gaussian measurements satisfies

$$m \geq C' \rho^6 \cdot w^2(\mathcal{D}(f, \mathbf{x}^\natural)).$$

How accurate is this result? Note that standard Gaussian measurements satisfy the assumptions of the corollary with ρ constant, and we need at least $w^2(\mathcal{D}(f, \mathbf{x}^\natural))$ standard normal measurements to recover the structured signal $\mathbf{x}^\natural$ with the complexity measure f . Therefore, the bound is correct up to the constant factor C' and the precise dependence on the eccentricity ρ .

2.6.4 Proof of Theorem 6.3: Setup and Step 1

To establish Theorem 6.3, we rely on Mendelson's Small Ball Method. The argument also depends on some deep ideas from the theory of generic chaining [32], but we only use these results in a naïve way.

Fix a cone K in $\mathbb{R}^d$ and define the set $E := K \cap \mathbf{S}^{d-1}$. Suppose that $\boldsymbol{\varphi}$ is a random vector in $\mathbb{R}^d$ that satisfies the conditions set out in Section 2.6.1 and construct an $m \times d$ random matrix Φ whose rows are independent copies of $\boldsymbol{\varphi}$. Proposition 5.1 implies that

$$\lambda_{\min}(\Phi; K) \geq \xi \sqrt{m} Q_{2\xi}(E) - 2W_m(E) - \xi t \quad \text{with probability} \geq 1 - e^{-t^2/2}. \quad (2.28)$$

This result holds for all $\xi > 0$ and $t > 0$. To establish Theorem 6.3, we must develop a constant lower bound for the marginal tail function $Q_{2\xi}(E)$, and we also need to compare the mean empirical width $W_m(E)$ with the conic Gaussian width $w(K)$.

2.6.5 Step 2: The marginal tail function

We begin with the lower bound for the marginal tail function $Q_{2\xi}$. This result is an easy consequence of the second moment method, also known as the Paley–Zygmund inequality. Let $\mathbf{u}$ be any vector in E . One version of the second moment method states that

$$\mathbb{P}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq 2\xi\} \geq \frac{[\mathbb{E} |\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| - 2\xi]_+^2}{\mathbb{E} |\langle \boldsymbol{\varphi}, \mathbf{u} \rangle|^2}. \quad (2.29)$$

To control the denominator on the right-hand side of (2.29), we use the sub-Gaussian marginal condition to estimate that

$$\mathbb{E} |\langle \boldsymbol{\varphi}, \mathbf{u} \rangle|^2 = \int_0^\infty 2s \cdot \mathbb{P}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq s\} ds \leq 4\sigma^2.$$

To bound the numerator on the right-hand side of (2.29), we use the nondegeneracy assumption: $\mathbb{E} |\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq \alpha$. Combining these results and taking the infimum over $\mathbf{u} \in E$, we reach

$$Q_{2\xi}(E) = \inf_{\mathbf{u} \in E} \mathbb{P}\{|\langle \boldsymbol{\varphi}, \mathbf{u} \rangle| \geq 2\xi\} \geq \frac{(\alpha - 2\xi)^2}{4\sigma^2} \quad (2.30)$$

for any ξ that satisfies $2\xi < \alpha$.

2.6.6 Step 3: The mean empirical width

Next, we demonstrate that the empirical width $W_m(E)$ is controlled by the conic Gaussian width $w(K)$. This argument requires sophisticated results from the theory of generic chaining [32]. First, observe that the vector $\mathbf{h} = m^{-1/2} \sum_{i=1}^m \varepsilon_i \boldsymbol{\varphi}_i$ inherits sub-Gaussian marginals from the centered sub-Gaussian distribution $\boldsymbol{\varphi}$. Indeed,

$$\mathbb{P}\{|\langle \mathbf{h}, \mathbf{u} \rangle| \geq t\} \leq C_1 e^{-c_1 t^2 / \sigma^2} \quad \text{for each } \mathbf{u} \in \mathbf{S}^{d-1}.$$

See [36, Sec. 5.2.3] for an introduction to sub-Gaussian random variables. In particular, we have the bound

$$\mathbb{P}\{|\langle \mathbf{h}, \mathbf{u} - \mathbf{v} \rangle| \geq t\} \leq C_1 e^{-c_1 t^2 / (\sigma^2 \|\mathbf{u} - \mathbf{v}\|^2)} \quad \text{for all } \mathbf{u}, \mathbf{v} \in \mathbb{R}^d.$$

Under the latter condition, the generic chaining theorem [32, Thm. 1.2.6] asserts that

$$W_m(E) = \mathbb{E} \sup_{\mathbf{u} \in E} \langle \mathbf{h}, \mathbf{u} \rangle \leq C_2 \sigma \cdot \gamma_2(E, \ell_2)$$

where γ_2 is a geometric functional. The precise definition of γ_2 is not important for our purposes because the majorizing measure theorem [32, Thm. 2.1.1] states that

$$\gamma_2(E, \ell_2) \leq C_3 \cdot \mathbb{E} \sup_{\mathbf{u} \in E} \langle \mathbf{g}, \mathbf{u} \rangle$$

where $\mathbf{g} \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$. It follows that

$$W_m(E) \leq C_4 \sigma \cdot \mathbb{E} \sup_{\mathbf{u} \in E} \langle \mathbf{g}, \mathbf{u} \rangle = C_4 \sigma \cdot w(K). \quad (2.31)$$

We have recalled that $E = K \cap \mathbf{S}^{d-1}$ to identify the conic Gaussian width $w(K)$.

2.6.7 Combining the bounds

Combine the bounds (2.28), (2.30), and (2.31) to discover that

$$\lambda_{\min}(\Phi; K) \geq \xi \sqrt{m} \cdot \frac{(\alpha - 2\xi)^2}{4\sigma^2} - 2C_4\sigma w(K) - \xi t \quad \text{with probability } \geq 1 - e^{-t^2/2},$$

provided that $2\xi < \alpha$. Select $\xi = \alpha/6$ to see that

$$\lambda_{\min}(\Phi; K) \geq \frac{1}{54} \cdot \frac{\alpha^3}{\sigma^2} \sqrt{m} - C_5\sigma w(K) - \frac{\alpha}{6}t \quad \text{with probability } \geq 1 - e^{-t^2/2}. \quad (2.32)$$

Using the eccentricity $\rho = \sigma/\alpha$, we simplify expression (2.32) to reach a bound for the minimum conic singular value of a sub-Gaussian random matrix Φ that satisfies the conditions set out in Section 2.6.1. This completes the proof of Theorem 6.3.

2.7 The bowling scheme

As we have seen in Theorem 6.3, sub-Gaussian sampling models exhibit behavior similar to the standard Gaussian measurement model. Yet there are many interesting problems where the random sampling matrix does not conform to the sub-Gaussian assumption. In this section, we explain how to adapt Mendelson's Small Ball Method to a range of other sampling ensembles. The key idea is to use the conic duality arguments from Section 2.4 to complete the estimate for the mean empirical width.

2.7.1 The mean empirical width of a descent cone

Let us state a simple duality result for the mean empirical width of a descent cone. This bound is based on the same principles as Proposition 4.5.

Proposition 7.1 (The mean empirical width of a descent cone). *Let $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ be a proper convex function, and fix a point $\mathbf{x} \in \mathbb{R}^d$. Assume that the subdifferential $\partial f(\mathbf{x})$ is nonempty and does not contain the origin. For any random vector $\boldsymbol{\varphi} \in \mathbb{R}^d$,*

$$W_m(\mathcal{D}(f, \mathbf{x}) \cap \mathbf{S}^{d-1}; \boldsymbol{\varphi}) \leq \mathbb{E} \inf_{\tau \geq 0} \text{dist}^2(\mathbf{h}, \tau \cdot \partial f(\mathbf{x})) \quad \text{where} \quad \mathbf{h} := \frac{1}{\sqrt{m}} \sum_{i=1}^m \varepsilon_i \boldsymbol{\varphi}_i.$$

The mean empirical width W_m is defined in (2.24). The random vectors $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_m$ are independent copies of $\boldsymbol{\varphi}$, and $\varepsilon_1, \dots, \varepsilon_m$ are independent Rademacher random variables.

Proof. The argument is identical with the proof of Proposition 4.5 once we replace the Gaussian vector $\mathbf{g}$ with the random vector $\mathbf{h}$. $\square$

2.7.2 The bowling scheme

We are now prepared to describe a general approach for convex signal recovery from independent random measurements.

The setup is similar to previous sections. Consider an unknown structured signal $\mathbf{x}^\natural \in \mathbb{R}^d$ and a complexity measure $f : \mathbb{R}^d \rightarrow \overline{\mathbb{R}}$ that is proper and convex. Let Φ be a known $m \times d$ sampling matrix, and suppose that we acquire m noisy linear measurements of the form $\mathbf{y} = \Phi \mathbf{x}^\natural + \mathbf{e}$. We wish to analyze the performance of the convex recovery method (2.2). Proposition 2.6 shows that we can accomplish this goal by finding a lower bound for the minimum conic singular value of the descent cone:

$$\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{x}^\natural)) \geq \text{???}. \quad (2.33)$$

We want to produce a bound of the form (2.33) when the rows of the measurement matrix Φ are independent copies of a random vector $\boldsymbol{\varphi}$. This problem falls within the scope of Mendelson's Small Ball Method. Introduce the index set $E := \mathcal{D}(f, \mathbf{x}^\natural) \cap \mathbb{S}^{d-1}$. In light of (2.23),

$$\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{x}^\natural)) = \inf_{\mathbf{u} \in E} \left(\sum_{i=1}^m |\langle \boldsymbol{\varphi}_i, \mathbf{u} \rangle|^2 \right)^{1/2}.$$

We follow Mendelson's general strategy to control the minimum conic singular value, but we propose a specific technique for bounding the mean empirical width that exploits the structure of the index set E .

THE BOWLING SCHEME

- (1) Apply Proposition 5.1 to bound the minimum conic singular value $\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{x}^\natural))$ below in terms of the marginal tail function $Q_{2\xi}(E; \boldsymbol{\varphi})$ and the mean empirical width $W_m(E; \boldsymbol{\varphi})$. The index set $E := \mathcal{D}(f; \mathbf{x}^\natural) \cap \mathbb{S}^{d-1}$.
- (2) Bound the marginal tail function $Q_{2\xi}(E; \boldsymbol{\varphi})$ below using a Paley–Zygmund inequality.
- (3') Apply Proposition 7.1 to control the mean empirical width $W_m(E; \boldsymbol{\varphi})$.

In other words, Step (3) of Mendelson's framework has been specialized to Step (3').

We refer to this instance of Mendelson’s Small Ball Method as *the bowling scheme*. The name is chosen as a salute to David Gross’s *golfing scheme*. Whereas the golfing scheme is based on dual optimality conditions for the signal recovery problem (2.2), the bowling scheme is based on the primal optimality condition through Proposition 2.6. In the bowling scheme, duality enters only when we are ready to estimate the mean empirical width.

In our experience, this idea has been successful whenever we understand how to bound the conic Gaussian width of the descent cone. The main distinction is that the random vector $\boldsymbol{\varphi}$ may not share the rotational invariance of the standard Gaussian distribution.

2.8 Example: Phase retrieval

To demonstrate how the bowling scheme works, we consider the question of phase retrieval. In this problem, we collect linear samples of an unknown signal, but we are only able to observe their magnitudes. To reconstruct the original signal, we must resolve the uncertainty about the phases (or signs) of the measurements. There is a natural convex program that can achieve this goal, and the bowling scheme offers an easy way to analyze the number of measurements that are required.

2.8.1 Phase retrieval by convex optimization

In the phase retrieval problem, we wish to recover a signal $\mathbf{x}^\natural \in \mathbb{R}^d$ from a family of measurements of the form

$$y_i = |\langle \boldsymbol{\psi}_i, \mathbf{x}^\natural \rangle|^2 \quad \text{for } i = 1, 2, 3, \dots, m. \quad (2.34)$$

The sampling ensemble $\boldsymbol{\psi}_1, \dots, \boldsymbol{\psi}_m$ consists of known vectors in $\mathbb{R}^d$. For clarity of presentation, we do not consider the case where the samples are noisy or complex-valued.

Although the samples do not initially appear linear, we can apply a lifting method proposed by Balan et al. [2]. Observe that

$$|\langle \boldsymbol{\psi}, \mathbf{x} \rangle|^2 = \boldsymbol{\psi}^\top \mathbf{x} \cdot \mathbf{x}^\top \boldsymbol{\psi} = \text{trace}(\mathbf{x} \mathbf{x}^\top \cdot \boldsymbol{\psi} \boldsymbol{\psi}^\top).$$

In view of this expression, it is appropriate to introduce the rank-one positive-semidefinite matrices

$$\mathbf{X}^\natural = (\mathbf{x}^\natural)(\mathbf{x}^\natural)^\top \in \mathbb{R}^{d \times d} \quad \text{and} \quad \boldsymbol{\Psi}_i = \boldsymbol{\psi}_i \boldsymbol{\psi}_i^\top \in \mathbb{R}^{d \times d} \quad \text{for } i = 1, 2, 3, \dots, m. \quad (2.35)$$

Then we can express the samples y_i as *linear* functions of the matrix $\mathbf{X}^\natural$:

$$y_i = \text{trace}(\mathbf{X}^\natural \cdot \Psi_i) \quad \text{for } i = 1, 2, 3, \dots, m. \quad (2.36)$$

Expression (2.36) coincides with the measurement model (2.1) we have been considering.

We can use convex optimization to reconstruct the unknown matrix $\mathbf{X}^\natural$. It is natural to minimize the Schatten 1-norm to promote low rank, but we also want to enforce the fact that $\mathbf{X}^\natural$ is positive semidefinite [9]. To that end, we consider the convex program

$$\begin{aligned} & \underset{\mathbf{X} \in \mathbb{R}^{d \times d}}{\text{minimize}} && \text{trace}(\mathbf{X}) && \text{subject to} && \mathbf{X} \succeq \mathbf{0} && \text{and} && y_i = \text{trace}(\mathbf{X} \Psi_i) \\ & && && && && && \text{for each } i = 1, 2, 3, \dots, m. \end{aligned} \quad (2.37)$$

This formulation involves the lifted variables (2.35). We say that the optimization problem (2.37) *recovers* $\mathbf{x}^\natural$ if the matrix $\mathbf{X}^\natural$ is the unique minimizer. Indeed, in this case, we can reconstruct the original signal by factorizing the solution to the optimization problem.

Remark 8.1 (Citation for convex phase retrieval). Formulation (2.37) was developed by a working group at the meeting “Frames for the finite world: Sampling, coding and quantization,” which took place at the American Institute of Mathematics in Palo Alto in August 2008. Most of the recent literature attributes this idea incorrectly.

2.8.2 Phase retrieval from Gaussian measurements

Recently, researchers have started to consider phase retrieval problems with random data; see [5] for example. In the simplest instance, we choose each sampling vector ψ_i independently from the standard normal distribution on $\mathbb{R}^d$:

$$\psi_i \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d).$$

Then each sampling matrix $\Psi_i = \psi_i \psi_i^\dagger$ follows a Wishart distribution. These random matrices do not have sub-Gaussian marginals, so we cannot apply Corollary 6.4 to study the performance of the optimization problem (2.37). Nevertheless, we can make short work of the analysis by using the bowling scheme.

Theorem 8.2 (Phase retrieval from Gaussian measurements). *Let $\mathbf{x}^\natural$ be a signal in $\mathbb{R}^d$. Let $\psi_i \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$ be independent standard Gaussian vectors, and consider random measurements $y_i = |\langle \psi_i, \mathbf{x}^\natural \rangle|^2$ for $i = 1, 2, 3, \dots, m$. Assuming that $m \geq Cd$, the convex phase retrieval problem (2.37) recovers $\mathbf{x}^\natural$ with probability at least $1 - e^{-cm}$. The numbers c and C are positive absolute constants.*

The sampling complexity $m \geq Cd$ established in Theorem 8.2 is qualitatively optimal. Indeed, a dimension-counting argument shows that we need at least $m \geq d$ nonadaptive linear measurements to reconstruct a general vector in $\mathbf{x}^\natural \in \mathbb{R}^d$.

Remark 8.3 (Extensions). There are a number of obvious improvements to Theorem 8.2 that follow with a little more effort. For example, it is clear that the convex phase retrieval method is stable. The exceedingly high success probability also allows us to establish uniform results for all d -dimensional vectors by means of net arguments and union bounds. Furthermore, the Gaussian assumption is inessential; it is possible to establish similar theorems for other sampling distributions. We leave these refinements for the avid reader.

2.8.3 Proof of Theorem 8.2: Setup

Let us rewrite the optimization problem (2.37) in a form that is more conducive to our methods of analysis. First, introduce the inner product space $\mathbb{R}_{\text{sym}}^{d \times d}$ of $d \times d$ symmetric matrices, equipped with the trace inner product $\langle \mathbf{A}, \mathbf{B} \rangle := \text{trace}(\mathbf{A}\mathbf{B})$ and the Frobenius norm $\|\cdot\|_F$. Define the linear operator

$$\Phi : \mathbb{R}_{\text{sym}}^{d \times d} \rightarrow \mathbb{R}^m \quad \text{where} \quad [\Phi(\mathbf{X})]_i = \langle \Psi_i, \mathbf{X} \rangle \quad \text{for } i = 1, 2, 3, \dots, m.$$

Collect the measurements into a vector $\mathbf{y} = (y_1, \dots, y_m)^\top \in \mathbb{R}^m$ and observe that $\mathbf{y} = \Phi(\mathbf{X}^\natural)$ because of expression (2.36). Next, define the convex indicator function of the positive-semidefinite cone:

$$\iota : \mathbb{R}_{\text{sym}}^{d \times d} \rightarrow \overline{\mathbb{R}} \quad \text{where} \quad \iota(\mathbf{X}) = \begin{cases} 0, & \mathbf{X} \text{ is positive semidefinite} \\ +\infty, & \text{otherwise.} \end{cases}$$

Introduce the convex regularizer

$$f : \mathbb{R}_{\text{sym}}^{d \times d} \rightarrow \overline{\mathbb{R}} \quad \text{where} \quad f(\mathbf{X}) = \text{trace}(\mathbf{X}) + \iota(\mathbf{X}).$$

With this notation, we can write (2.37) in the form

$$\underset{\mathbf{X} \in \mathbb{R}_{\text{sym}}^{d \times d}}{\text{minimize}} \quad f(\mathbf{X}) \quad \text{subject to} \quad \mathbf{y} = \Phi(\mathbf{X}). \quad (2.38)$$

Formulation (2.38) matches our core problem (2.2) with the error vector $\mathbf{e} = \mathbf{0}$ and error tolerance $\eta = 0$.

Proposition 2.6 demonstrates that $\mathbf{X}^\natural$ is the unique solution of (2.38) whenever

$$\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{X}^\natural)) > 0.$$

We must determine how many measurements m suffice for this event to hold with high probability.

2.8.4 Step 1: The nonnegative empirical process bound

Define the set

$$E := \{U \in \mathcal{D}(f, X^{\natural}) : \|U\|_F = 1\} \subset \mathbb{R}_{\text{sym}}^{d \times d}.$$

Proposition 5.1 demonstrates that

$$\lambda_{\min}(\Phi; \mathcal{D}(f, X^{\natural})) = \inf_{U \in E} \left(\sum_{i=1}^m |\langle \Psi_i, U \rangle|^2 \right)^{1/2} \geq \xi \sqrt{m} Q_{2\xi}(E) - 2 W_m(E) - \xi t \quad (2.39)$$

with probability at least $1 - e^{-t^2/2}$. In this setting, the marginal tail function is defined as

$$Q_{2\xi}(E) := \inf_{U \in E} \mathbb{P}\{|\langle \Psi_1, U \rangle| \geq 2\xi\}.$$

The mean empirical width is defined as

$$W_m(E) := \mathbb{E} \sup_{U \in E} \langle H, U \rangle \quad \text{where} \quad H := \frac{1}{\sqrt{m}} \sum_{i=1}^m \varepsilon_i \Psi_i.$$

Here, $\{\varepsilon_i\}$ is an independent family of Rademacher random variables, independent of everything else.

2.8.5 Step 2: The marginal tail function

We can use the Paley–Zygmund inequality to show that

$$Q_1(E) = \inf_{U \in E} \mathbb{P}\{|\langle \Psi_1, U \rangle| \geq 1\} \geq c_0. \quad (2.40)$$

We have implicitly chosen $\xi = \frac{1}{2}$, and c_0 is a positive absolute constant.

2.8.5.1 The tail bound

To perform this estimate, we apply the Paley–Zygmund inequality in the form

$$\mathbb{P}\left\{|\langle \Psi_1, U \rangle|^2 \geq \frac{1}{2} (\mathbb{E} |\langle \Psi_1, U \rangle|^2)\right\} \geq \frac{1}{4} \cdot \frac{(\mathbb{E} |\langle \Psi_1, U \rangle|^2)^2}{\mathbb{E} |\langle \Psi_1, U \rangle|^4}.$$

The easiest way to treat the expectation in the denominator is to invoke Gaussian hypercontractivity [17, Sec. 3.2]. Indeed,

$$\left(\mathbb{E} |\langle \Psi_1, U \rangle|^4\right)^{1/4} \leq C_0 \left(\mathbb{E} |\langle \Psi_1, U \rangle|^2\right)^{1/2}$$

because $\langle \Psi_1, U \rangle$ is a second-order polynomial in the entries of ψ_1 . Combine the last two displays to obtain

$$\mathbb{P}\left\{|\langle \Psi_1, U \rangle|^2 \geq \frac{1}{2} \left(\mathbb{E} |\langle \Psi_1, U \rangle|^2\right)\right\} \geq \frac{1}{4 \cdot C_0^4} = c_0.$$

We can bound the remaining expectation by means of an explicit calculation. Assuming that $U \in E$,

$$\mathbb{E} |\langle \Psi_1, U \rangle|^2 = 3 \sum_{i=1}^m |u_{ii}|^2 + 2 \sum_{i,j=1}^m |u_{ij}|^2 + \left| \sum_{i=1}^m u_{ii} \right|^2 \geq 2.$$

We have used the fact that U is a symmetric matrix with unit Frobenius norm. In conclusion,

$$\mathbb{P}\left\{|\langle \Psi_1, U \rangle|^2 \geq 1\right\} \geq c_0 \quad \text{for each } U \in E.$$

This inequality implies (2.40).

2.8.6 Step 3': The mean empirical width of the descent cone

We can apply Proposition 7.1 to demonstrate that the mean empirical width satisfies

$$W_m(E) \leq C_1 \sqrt{d} \quad \text{for } m \geq C_2 d. \quad (2.41)$$

The numbers C_1 and C_2 are positive, absolute constants.

2.8.6.1 The width bound

The bound holds trivially when $X^\natural = \mathbf{0}$, so we may assume that the unknown matrix is nonzero. Select a coordinate system where

$$X^\natural = \begin{bmatrix} a & \mathbf{0}^\top \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}_{\text{sym}}^{d \times d} \quad \text{where } a > 0.$$

Recall that the matrix $\mathbf{H} = m^{-1/2} \sum_{i=1}^m \varepsilon_i \boldsymbol{\Psi}_i$, where $\boldsymbol{\Psi}_i = \boldsymbol{\psi}_i \boldsymbol{\psi}_i^\top$ and $\boldsymbol{\psi}_i \sim \text{NORMAL}(\mathbf{0}, \mathbf{I}_d)$. Partition $\mathbf{H}$ conformally with $\mathbf{X}^\natural$:

$$\mathbf{H} = \begin{bmatrix} h_{11} & \mathbf{h}_{21}^\top \\ \mathbf{h}_{21} & \mathbf{H}_{22} \end{bmatrix}.$$

Define the random parameter $\tau = \|\mathbf{H}_{22}\|$.

$$W_m(E) = \mathbb{E} \sup_{\mathbf{U} \in E} \langle \mathbf{H}, \mathbf{U} \rangle \leq \left(\mathbb{E} \text{dist}_{\mathbb{F}}^2(\mathbf{H}, \tau \cdot \partial f(\mathbf{X}^\natural)) \right)^{1/2}. \quad (2.42)$$

Using standard calculus rules for subdifferentials [28, Chap. 23], we determine that

$$\partial f(\mathbf{X}^\natural) = \left\{ \begin{bmatrix} 1 & \mathbf{0}^\top \\ \mathbf{0} & \mathbf{Y} \end{bmatrix} \in \mathbb{R}_{\text{sym}}^{d \times d} : \lambda_{\max}(\mathbf{Y}) \leq 1 \right\}.$$

We write $\lambda_{\max}$ denotes the maximum eigenvalue of a symmetric matrix. Proposition 7.1 delivers the width bound,

$$\mathbb{E} \text{dist}_{\mathbb{F}}^2(\mathbf{H}, \partial f(\mathbf{X}^\natural)) = \mathbb{E} (h_{11} - \tau)^2 + 2 \mathbb{E} \|\mathbf{h}_{21}\|^2 + \mathbb{E} \inf_{\lambda_{\max}(\mathbf{Y}) \leq 1} \|\mathbf{H}_{22} - \tau \cdot \mathbf{Y}\|_{\mathbb{F}}^2. \quad (2.43)$$

By construction, the third term on the right-hand side of (2.43) is zero. By direct calculation, the second term on the right-hand side of (2.43) satisfies

$$\mathbb{E} \|\mathbf{h}_{21}\|^2 = d - 1. \quad (2.44)$$

Finally, we turn to the first term on the right-hand side of (2.43). Relatively crude bounds suffice here. By interlacing of singular values,

$$\tau = \|\mathbf{H}_{22}\| \leq \|\mathbf{H}\| = \frac{1}{\sqrt{m}} \left\| \sum_{i=1}^m \varepsilon_i \boldsymbol{\psi}_i \boldsymbol{\psi}_i^\top \right\|.$$

Standard net arguments, such as those in [36, Sec. 5.4.1], demonstrate that

$$\mathbb{P}\{\|\mathbf{H}\| \geq C_3 \sqrt{d}\} \leq e^{-c_1 d}, \quad \text{provided that } m \geq C_2 d.$$

Together, the last two displays imply that $\mathbb{E} \tau^2 \leq C_4 d$. Therefore,

$$\mathbb{E} (h_{11} - \tau)^2 \leq C_5 d. \quad (2.45)$$

Introducing (2.43), (2.44), and (2.45) into (2.42), we arrive at the required bound (2.41).

Remark 8.4 (Other sampling distributions). The only challenging part of the calculation is the bound on $\|\mathbf{H}\|$. For more general sampling distributions, we can easily obtain the required estimate from the matrix moment inequality [7, Thm. A.1].

2.8.7 Combining the bounds

Assume that $m \geq C_2 d$. Combine the estimates (2.39), (2.40), and (2.41) to reach

$$\lambda_{\min}(\Phi; \mathcal{D}(f, \mathbf{X}^\natural)) \geq c_2 \sqrt{m} - C_6 \sqrt{d} - \frac{1}{2}t$$

with probability at least $1 - e^{-t^2/2}$. Choosing $t = c_3 \sqrt{m}$, we find that the minimum conic singular value is positive with probability at least $1 - e^{-c_4 m}$. In this event, Proposition 2.6 implies that $\mathbf{X}^\natural$ is the unique solution to the phase retrieval problem (2.37). This observation completes the proof of Theorem 8.2.

Acknowledgements JAT gratefully acknowledges support from ONR award N00014-11-1002, AFOSR award FA9550-09-1-0643, and a Sloan Research Fellowship. Thanks are also due to the Moore Foundation.

References

1. D. Amelunxen, M. Lotz, M.B. McCoy, J.A. Tropp, Living on the edge: phase transitions in convex programs with random data. *Inf. Inference* **3**(3), 224–294 (2014). Available at <http://arXiv.org/abs/1303.6672>
2. R. Balan, B.G. Bodmann, P.G. Casazza, D. Edidin, Painless reconstruction from magnitudes of frame coefficients. *J. Fourier Anal. Appl.* **15**(4), 488–501 (2009)
3. S. Boucheron, G. Lugosi, P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence* (Oxford University Press, Oxford, 2013)
4. T.T. Cai, T. Liang, A. Rakhlin, Geometrizing local rates of convergence for linear inverse problems (April 2014). Available at <http://arXiv.org/abs/1404.4408>
5. E.J. Candès, T. Strohmer, V. Voroninski, PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming. *Commun. Pure Appl. Math.* **66**(8), 1241–1274 (2013)
6. V. Chandrasekaran, B. Recht, P.A. Parrilo, A.S. Willsky, The convex geometry of linear inverse problems. *Found. Comput. Math.* **12**(6), 805–849 (2012)
7. R.Y. Chen, A. Gittens, J.A. Tropp, The masked sample covariance estimator: an analysis via the matrix Laplace transform method. *Inf. Inference* **1**, 2–20 (2012)
8. K.R. Davidson, S.J. Szarek, Local operator theory, random matrices and Banach spaces, in *Handbook of the Geometry of Banach Spaces*, vol. I (North-Holland, Amsterdam, 2001), pp. 317–366
9. M. Fazel. *Matrix Rank Minimization with Applications*. Ph.D. thesis, Stanford University, 2002.

10. R. Foygel, L. Mackey, Corrupted sensing: novel guarantees for separating structured signals. *Trans. Inf. Theory* **60**(2), 1223–1247 (2014)
11. Y. Gordon, Some inequalities for Gaussian processes and applications. *Isr. J. Math.* **50**(4), 265–289 (1985)
12. Y. Gordon, On Milman’s inequality and random subspaces which escape through a mesh in $\mathbf{R}^n$, in *Geometric Aspects of Functional Analysis (1986/1987)*. Lecture Notes in Mathematics, vol. 1317 (Springer, Berlin, 1988), pp. 84–106
13. D. Gross, Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inf. Theory* **57**(3), 1548–1566 (2011)
14. V. Koltchinskii, S. Mendelson, Bounding the smallest singular value of a random matrix without concentration (December 2013). Available at <http://arXiv.org/abs/1312.3580>
15. R. Latała, K. Oleszkiewicz, On the best constant in the Khinchin-Kahane inequality. *Studia Math.* **109**(1), 101–104 (1994)
16. G. Lecué, S. Mendelson, Compressed sensing under weak moment assumptions (January 2014). Available at <http://arXiv.org/abs/1401.2188>
17. M. Ledoux, M. Talagrand, *Probability in Banach Spaces: Isoperimetry and Processes* (Springer, Berlin, 1991)
18. S. Mendelson, Empirical processes with a bounded ψ_1 diameter. *Geom. Funct. Anal.* **20**(4), 988–1027 (2010)
19. S. Mendelson, A remark on the diameter of random sections of convex bodies (December 2013). Available at <http://arXiv.org/abs/1312.3608>
20. S. Mendelson, Learning without concentration. *J. Assoc. Comput. Mach.* (2014, to appear). **62**(3), (2015). Available at <http://arXiv.org/abs/1401.0304>
21. S. Mendelson, Learning without concentration for general loss functions (October 2014). Available at <http://arXiv.org/abs/1410.3192>
22. S. Mendelson, A. Pajor, N. Tomczak-Jaegermann, Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geom. Funct. Anal.* **17**(4), 1248–1282 (2007)
23. V. Milman, G. Schechtman, *Asymptotic Theory of Finite-Dimensional Normed Linear Spaces*. Number 1200 in LNM (Springer, New York, 1986)
24. S. Oymak, B. Hassibi, New null space results and recovery thresholds for matrix rank minimization. Partial results presented at ISIT 2011 (2010). Available at <http://arXiv.org/abs/1011.6326>
25. S. Oymak, B. Hassibi, Sharp MSE bounds for proximal denoising. Partial results presented at Allerton 2012 (March 2013). Available at <http://arxiv.org/abs/1305.2714>
26. S. Oymak, C. Thrampoulides, B. Hassibi, Simple bounds for noisy linear inverse problems with exact side information (December 2013). Available at <http://arXiv.org/abs/1312.0641>
27. G. Pisier, *The Volume of Convex Bodies and Banach Space Geometry* (Cambridge University Press, Cambridge, 1989)
28. R.T. Rockafellar, *Convex Analysis* (Princeton University Press, Princeton, 1970)
29. M. Rudelson, R. Vershynin, On sparse reconstruction from Fourier and Gaussian measurements. *Commun. Pure Appl. Math.* **61**(8), 1025–1045 (2008)
30. M. Sion, On general minimax theorems. *Pac. J. Math.* **8**, 171–176 (1958)
31. M. Stojnic, Various thresholds for ℓ_1 -optimization in compressed sensing (2009). Available at <http://arXiv.org/abs/0907.3666>
32. M. Talagrand, *The Generic Chaining. Upper and Lower Bounds of Stochastic Processes*. Springer Monographs in Mathematics (Springer, Berlin, 2005)
33. C. Thrampoulides, S. Oymak, B. Hassibi, Simple error bounds for regularized noisy linear inverse problems. Appeared at ISIT 2014 (January 2014). Available at <http://arXiv.org/abs/1401.6578>
34. J. Tropp, M. Wakin, M. Duarte, D. Baron, R. Baraniuk, Random filters for compressive sampling and reconstruction, in *Proceedings of the 2006 IEEE International Conference on Acoustics, Speech and Signal Processing, 2006 (ICASSP 2006)*, vol. 3, May 2006, pp. III–875

- 35. A.W. van der Vaart, J.A. Wellner, *Weak Convergence and Empirical Processes*. Springer Series in Statistics (Springer, New York, 1996). With applications to statistics.
- 36. R. Vershynin, Introduction to the non-asymptotic analysis of random matrices, in *Compressed Sensing* (Cambridge University Press, Cambridge, 2012), pp. 210–268
- 37. G.A. Watson, Characterization of the subdifferential of some matrix norms. *Linear Algebra Appl.* **170**, 33–45 (1992)

Sampling Theory, a Renaissance
Compressive Sensing and Other Developments

Pfander, G.E. (Ed.)

2015, XIV, 532 p. 69 illus., 40 illus. in color., Hardcover

ISBN: 978-3-319-19748-7

A product of Birkhäuser Basel