Rice University

Random Observations on Random Observations: Sparse Signal Acquisition and Processing

by

Mark A. Davenport

A Thesis Submitted
in Partial Fulfillment of the
Requirements for the Degree

Doctor of Philosophy

Approved, Thesis Committee:

Richard G. Baraniuk, Chair,
Victor E. Cameron Professor,
Electrical and Computer Engineering

Kevin F. Kelly, Associate Professor,
Electrical and Computer Engineering

Mark P. Embree, Professor,
Computational and Applied Mathematics

Piotr Indyk, Associate Professor with Tenure,
Electrical Engineering and Computer Science,
Massachusetts Institute of Technology

Ronald A. DeVore, Walter E. Koss Professor,
Mathematics, Texas A&M University

Houston, Texas

August 2010

Abstract

Random Observations on Random Observations:
Sparse Signal Acquisition and Processing

by

Mark A. Davenport

In recent years, signal processing has come under mounting pressure to accommodate the increasingly high-dimensional raw data generated by modern sensing systems. Despite extraordinary advances in computational power, processing the signals produced in application areas such as imaging, video, remote surveillance, spectroscopy, and genomic data analysis continues to pose a tremendous challenge. Fortunately, in many cases these high-dimensional signals contain relatively little information compared to their ambient dimensionality. For example, signals can often be well-approximated as a sparse linear combination of elements from a known basis or dictionary.

Traditionally, sparse models have been exploited only after acquisition, typically for tasks such as compression. Recently, however, the applications of sparsity have greatly expanded with the emergence of compressive sensing, a new approach to data acquisition that directly exploits sparsity in order to acquire analog signals more efficiently via a small set of more general, often randomized, linear measurements. If properly chosen, the number of measurements can be much smaller than the number of Nyquist-rate samples. A common theme in this research is the use of randomness in signal acquisition, inspiring the design of hardware systems that directly implement random measurement protocols.

This thesis builds on the field of compressive sensing and illustrates how sparsity can be exploited to design efficient signal processing algorithms at all stages of the

information processing pipeline, with a particular focus on the manner in which randomness can be exploited to design new kinds of acquisition systems for sparse signals. Our key contributions include: *(i)* exploration and analysis of the appropriate properties for a sparse signal acquisition system; *(ii)* insight into the useful properties of random measurement schemes; *(iii)* analysis of an important family of algorithms for recovering sparse signals from random measurements; *(iv)* exploration of the impact of noise, both structured and unstructured, in the context of random measurements; and *(v)* algorithms that process random measurements to directly extract higher-level information or solve inference problems without resorting to full-scale signal recovery, reducing both the cost of signal acquisition and the complexity of the post-acquisition processing.

Acknowledgements

One of the best parts of graduate school at Rice has been the chance to work with so many great people. I have been incredibly fortunate to work with an amazing group of collaborators: Rich Baraniuk, Petros Boufounos, Ron DeVore, Marco Duarte, Chin Hegde, Kevin Kelly, Jason Laska, Stephen Schnelle, Clay Scott, J.P. Slavinsky, Ting Sun, Dharmpal Takhar, John Treichler, and Mike Wakin. Much of this thesis is built upon work with these collaborators, as I have noted in each chapter. Thank you all for your help.

I would also like to thank my committee for their valuable suggestions and contributions: to Kevin Kelly for providing a constant stream of off-the-wall ideas (a single-pixel camera?); to Mark Embree both for his terrific courses on numerical and functional analysis and for taking the time to listen to and provide feedback on so many ideas over the years; to Piotr Indyk for providing a completely new perspective on my research and for hosting me at MIT and visiting us at Rice; to Ron DeVore for pouring so much energy into de-mystifying compressive sensing during his year at Rice and for remaining such an encouraging collaborator; and most of all to my advisor Rich Baraniuk for somehow pushing me just hard enough over the years. I would not be the researcher I am today were it not for the countless hours Rich spent helping me through tough problems in his office, over coffee, or during home improvement projects. Rich is the definition of enthusiasm and has been a true inspiration.

I could not have asked for a more energetic and fun group of people to work

with over the years. Thanks to Albert, Anna, Aswin, Chris, Clay, Dharmpal, Drew, Eva, Gareth, Ilan, Jarvis, Jason, J.P., Joel, John, Jyoti, Kadim, Khanh, Kevin, Kia, Laurent, Liz, Manjari, Marco, Mark, Martin, Matthew, Mike, Mona, Neelsh, Peter, Petros, Piotr, Prashant, Ray, Rich, Richard, Robin, Ron, Sarah, Sanda, Shri, Sid, Sidney, Stephen, Ting, Wotao, and many more. You made Duncan Hall a great place to work, and also a great place to avoid doing work. I will always remember all the time I spent we spent at Valhalla, on "POTUS of the day", talking about Lost, eating J.P.'s doughnuts, and occasionally getting some work done too. I would also like to thank Chris, Jason, and Mike for providing the most rewarding distraction of them all: *Rejecta Mathematica*.

Finally, I must thank my wonderful friends and family for helping me to keep things in perspective over the years, especially: my friends Brandon, and Elvia; Kim's family Bob, Lan, Lynh, and Robert; my brothers Blake and Brian and their families; my grandparents Joe, Margaret, Cliff, and LaVonne; and my parents, Betty and Dennis. Most importantly, I am deeply grateful to my amazing wife Kim for her unending love, support, and encouragement.

Contents

List of Illustrations

List of Algorithms

Dedicated to my wife Kim.

Without your constant support

I would still be stuck on chapter one.

Chapter 1

Introduction

1.1 Models in Signal Processing

At its core, signal processing is concerned with efficient algorithms for acquiring and extracting information from signals or data. In order to design such algorithms for a particular problem, we must have accurate *models* for the signals of interest. These can take the form of generative models, deterministic classes, or probabilistic Bayesian models. In general, models are useful for incorporating *a priori* knowledge to help distinguish classes of interesting or probable signals from uninteresting or improbable signals, which can help us to efficiently and accurately acquire, process, compress, and communicate data and information.

For much of its history, signal processing has focused on signals produced by physical systems. Many natural and manmade systems can be modeled as linear systems, thus, it is natural to consider signal models that complement this kind of linear structure. This notion has been incorporated into modern signal processing by modeling signals as *vectors* living in an appropriate *vector space*. This captures the linear structure that we often desire, namely that if we add two signals together we obtain a new, physically meaningful signal. Moreover, vector spaces allow us to apply

intuitions and tools from geometry in $\mathbb{R}^3$, such as lengths, distances, and angles, to describe and compare our signals of interest. This is useful even when our signals live in high-dimensional or infinite-dimensional spaces.

Such linear models are widely applicable and have been studied for many years. For example, the theoretical foundation of *digital signal processing* (DSP) is the pioneering work of Whittaker, Nyquist, Kotelnikov, and Shannon [1–4] on the sampling of continuous-time signals. Their results demonstrate that bandlimited, continuous-time signals, which define a vector space, can be exactly recovered from a set of uniformly-spaced samples taken at the Nyquist rate of twice the bandlimit. Capitalizing on this discovery, signal processing has moved from the analog to the digital domain and ridden the wave of Moore's law. Digitization has enabled the creation of sensing and processing systems that are more robust, flexible, cheaper and, therefore, more widely-used than their analog counterparts.

As a result of this success, the amount of data generated by sensing systems has grown from a trickle to a torrent. Unfortunately, in many important and emerging applications, the resulting Nyquist rate is so high that we end up with too many samples, at which point many algorithms become overwhelmed by the so-called "curse of dimensionality" [5]. Alternatively, it may simply be too costly, or even physically impossible, to build devices capable of acquiring samples at the necessary rate [6, 7]. Thus, despite extraordinary advances in computational power, acquiring and processing signals in application areas such as imaging, video, remote surveillance, spectroscopy, and genomic data analysis continues to pose a tremendous challenge. Moreover, simple linear models often fail to capture much of the structure present in such signals.

In response to these challenges, there has been a surge of interest in recent years across many fields in a variety of *low-dimensional signal models*. Low-dimensional models provide a mathematical framework for capturing the fact that in many cases

(a) (b)

Figure 1.1: Sparse representation of an image via a multiscale wavelet transform. (a) Original image (b) Wavelet representation. Large coefficients are represented by light pixels, while small coefficients are represented by dark pixels. Observe that most of the wavelet coefficients are near zero.

these high-dimensional signals contain relatively little information compared to their ambient dimensionality. For example, signals can often be well-approximated as a linear combination of just a few elements from a known basis or dictionary, in which case we say that the signal is *sparse*. Sparsity has been exploited heavily in fields such as image processing for tasks like compression and denoising [8], since the multiscale wavelet transform [9] provides nearly sparse representations for natural images. An example is shown in Figure 1.1. Sparsity also figures prominently in the theory of statistical estimation and model selection [10] and in the study of the human visual system [11].

Sparsity is a highly nonlinear model, since the choice of which dictionary elements are used can change from signal to signal [12]. In fact, it is easy to show that the set of all sparse signals consists of not one subspace but the union of a combinatorial number of subspaces. As a result, we must turn to nonlinear algorithms in order to exploit sparse models. This nonlinear nature has historically limited the use of sparse models due to the apparent need for computationally complex algorithms in order to

exploit sparsity. In recent years, however, there has been tremendous progress in the design of efficient algorithms that exploit sparsity. In particular, sparsity lies at the heart of the emerging field of *compressive signal processing* (CSP).

1.2 Compressive Signal Processing

1.2.1 Compressive sensing and signal acquisition

The Nyquist-Shannon sampling theorem states that a certain minimum amount of sampling is required in order to perfectly capture an arbitrary bandlimited signal. On the other hand, if our signal is sparse in a known basis, we can vastly reduce how many numbers must be stored, far below the supposedly minimal number of required samples. This suggests that for the case of sparse signals, we might be able to do better than classical results would suggest. This is the fundamental idea behind the emerging field of *compressive sensing* (CS) [13–19].

While this idea has only recently gained significant traction in the signal processing community, there have been hints in this direction dating back as far as 1795 with the work of Prony on the estimation of the parameters associated with a small number of complex exponentials sampled in the presence of noise [20]. The next theoretical leap came in the early 1900's, when Carathéodory showed that a positive linear combination of *any* K sinusoids is uniquely determined by its value at $t = 0$ and at *any* other $2K$ points in time [21, 22]. This represents far fewer samples than the number of Nyquist-rate samples when K is small and the range of possible frequencies is large. We then fast-forward to the 1990's, when this work was generalized by Feng, Bresler, and Venkataramani, who proposed a practical sampling scheme for acquiring signals consisting of K components with nonzero bandwidth (as opposed to pure sinusoids), reaching somewhat similar conclusions [23–27]. Finally, in the early 2000's Vetterli, Marziliano, and Blu proposed a sampling scheme for certain classes of

non-bandlimited signals that are governed by only K parameters, showing that these signals can be sampled and recovered from only $2K$ samples [28].

In a somewhat different setting, Beurling considered the problem of when we can recover a signal by observing only a piece of its Fourier transform. He proposed a method to extrapolate from these observations to determine the entire Fourier transform [29]. One can show that if the signal consists of a finite number of impulses, then Beurling's approach will correctly recover the entire Fourier transform (of this non-bandlimited signal) from *any* sufficiently large piece of its Fourier transform. His approach — to find the signal with smallest ℓ_1 norm among all signals agreeing with the acquired Fourier measurements — bears a remarkable resemblance to some of the algorithms used in CS.

Building on these results, CS has emerged as a new framework for signal acquisition and sensor design that enables a potentially large reduction in the cost of acquiring signals that have a sparse or compressible representation. CS builds on the work of Candès, Romberg, Tao [13–17], and Donoho [18], who showed that a signal having a sparse representation can be recovered *exactly* from a small set of linear, nonadaptive *compressive measurements*. However, CS differs from classical sampling is two important respects. First, rather than sampling the signal at specific points in time, CS systems typically acquire measurements in the form of inner products between the signal and more general test functions. We will see in this thesis that *randomness* often plays a key role in the design of these test functions. Secondly, the two frameworks differ in the manner in which they deal with *signal recovery*, i.e., the problem of recovering the original signal from the compressive measurements. In the Nyquist-Shannon framework, signal recovery is achieved through sinc interpolation — a linear process that requires little computation and has a simple interpretation. In CS, however, signal recovery is achieved using nonlinear and relatively expensive optimization-based or iterative algorithms [30–45]. See [46] for an overview of these

methods.

CS has already had notable impacts on medical imaging [47–50]. In one study it has been demonstrated to enable a speedup by a factor of seven in pediatric MRI while preserving diagnostic quality [51]. Moreover, the broad applicability of this framework has inspired research that extends the CS framework by proposing practical implementations for numerous applications, including sub-Nyquist sampling systems [52–55], compressive imaging architectures [56–58], and compressive sensor networks [59, 60].

1.2.2 Compressive domain processing

Despite the intense focus of the CS community on the problem of signal recovery, it is not actually necessary in many signal processing applications. In fact, most of the field of digital signal processing (DSP) is actually concerned with solving *inference* problems, i.e., extracting only certain information from measurements. For example, we might aim to detect the presence of a signal of interest, classify among a set of possible candidate signals, estimate some function of the signal, or filter out a signal that is not of interest before further processing. While one could always attempt to recover the full signal from the compressive measurements and then solve such problems using traditional DSP techniques, this approach is typically suboptimal in terms of both accuracy and efficiency.

This thesis takes some initial steps towards a general framework for what we call *compressive signal processing* (CSP), an alternative approach in which signal processing problems are solved directly in the compressive measurement domain *without* first resorting to a full-scale signal reconstruction. This can take on many meanings. For example, in [61] sparsity is leveraged to perform classification with very few random measurements, and a variety of additional approaches to detection, classification, and estimation from compressive measurements are further examined in this thesis,

along with an approach to filtering compressive measurements to remove interference. A general theme of these efforts is that compressive measurements are *information scalable* — complex inference tasks like recovery require many measurements, while comparatively simple tasks like detection require far fewer measurements.

While this work builds on the CS framework, it also shares a close relationship with the field of *data streaming* algorithms, which is concerned with processing large streams of data using efficient algorithms. The data streaming community has examined a huge variety of problems over the past several years. In the data stream setting, one is typically interested in estimating some function of the data stream (such as an ℓ_p norm, a histogram, or a linear functional) based on a linear "sketch". For a concise review of these results see [62], or see [63] for a more recent overview of data streaming algorithms in the context of CS. The results from this community also demonstrate that in many cases it is possible to save in terms of both the required number of measurements as well as the required amount of computation if we directly solve the problem of interest without resorting to recovering the original signal. Note, however, that while data streaming algorithms typically design a sketch to target a specific problem of interest, the CSP approach is to use the same generic compressive measurements to solve a wide range of potential inference problems.

1.3 Overview and Contributions

This thesis builds on the field of compressive sensing and illustrates how sparsity can be exploited to design efficient signal processing algorithms at all stages of the information processing pipeline, with a particular focus on the manner in which randomness can be exploited to design new kinds of acquisition systems for sparse signals. Our key contributions include:

- exploration and analysis of the appropriate properties for a sparse signal acqui-

sition system;

- insight into the useful properties of random measurement schemes;
- analysis of an important family of algorithms for recovering sparse signals from random measurements;
- exploration of the impact of noise, both structured and unstructured, in the context of random measurements; and
- algorithms that process random measurements to directly extract higher-level information or solve inference problems without resorting to full-scale signal recovery, both reducing the cost of signal acquisition and reducing the complexity of the post-acquisition processing.

For clarity, these contributions are organized into four parts.

In **Part I** we introduce the concept of sparse signal models. After a brief discussion of mathematical preliminaries and notation, **Chapter 2** provides a review of sparse and compressible models. Additionally, we give an overview of the sparse approximation algorithms that will play a crucial role in CS.

Next, in **Part II** we describe methods for sparse signal acquisition. We begin this discussion in **Chapter 3** by exploring the properties that we will require our signal acquisition system to satisfy to ensure that we preserve the information content of sparse signals. This will lead us to the notion of stable embeddings and the *restricted isometry property* (RIP). We will explore this property, providing an argument for its necessity when dealing with certain kinds of noise, and providing a brief overview of the theoretical implications of the RIP in CS. We will also establish lower bounds on how many measurements are required for a matrix to satisfy the RIP.

Chapter 4 then describes an argument that certain random matrices will satisfy the RIP. We begin with an overview of sub-Gaussian distributions — a family of

probability distributions which behave like Gaussian distributions in a certain respect in high dimensions. Specifically, we prove that sub-Gaussian distributions exhibit a concentration of measure property, and then we exploit this property to argue that when a sub-Gaussian matrix has sufficiently many rows, it will satisfy the RIP with high probability. We also provide some discussion on the role of randomness and probabilistic guarantees within the broader field of CS, and describe how these techniques can also be extended to models beyond sparsity.

In **Chapter 5** we discuss various strategies for implementing these kinds of measurement techniques in systems for acquiring real-world signals. We primarily focus on two signal acquisition architectures: the single-pixel camera and the random demodulator. The single-pixel camera uses a Texas Instruments DMD array and a single light sensor to optically compute inner products between an image and random patterns. By changing these patterns over time, we can build up a collection of random measurements of an image. The random demodulator provides a CS-inspired hardware architecture for acquiring wideband analog signals. In both cases, we demonstrate that we can adapt the finite-dimensional acquisition framework described in the previous chapters to acquire continuous-time, analog signals.

Part III shifts the focus to the problem of recovering sparse signals from the kind of measurements produced by the systems described in Part II. **Chapter 6** begins by providing an RIP-based theoretical framework for analyzing orthogonal greedy algorithms. First, we provide an RIP-based analysis of the classical algorithm of *Orthogonal Matching Pursuit* (OMP) when applied to recovering sparse signals in the noise-free setting. We show that in this setting, if our measurement system satisfies the RIP, then OMP will succeed in recovering a K-sparse signal in exactly K iterations. We then extend this analysis and use the same techniques to establish a simple proof that under even weaker assumptions, *Regularized OMP* (ROMP) will also succeed in recovering K-sparse signals.

Chapter 7 then analyzes the potential impact of noise on the acquisition and recovery processes. We first discuss the case where noise is added to the measurements, and examine the performance of an oracle-assisted recovery algorithm. We conclude that the performance of most standard sparse recovery algorithms is near-optimal in that it matches the performance of an oracle-assisted algorithm. Moreover, in this setting the impact of the noise is well-controlled in the sense that the resulting recovery error is comparable to the size of the measurement noise. We then consider the case where noise is added to the signal itself. In the case of white noise we show that compressive measurement systems will amplify this noise by an amount determined only by the number of measurements taken. Specifically, we observe that the recovered signal-to-noise ratio will decrease by 3dB each time the number of measurements is reduced by a factor of 2. This suggests that in low signal-to-noise ratio (SNR) settings, CS-based acquisition systems will be highly susceptible to noise.

In **Chapter 8** we consider the impact of more structured noise. Specifically, we analyze the case where the noise itself is also sparse. We demonstrate that in addition to satisfying the RIP, the same random matrices considered in Chapter 4 satisfy an additional property that leads to measurements that are guaranteed to be robust to a small number of arbitrary corruptions and to other forms of sparse measurement noise. We propose an algorithm dubbed *Justice Pursuit* that can exploit this structure to recover sparse signals in the presence of corruption. We then show that this structure can be viewed as an example of a more general phenomenon. Specifically, we propose a definition of *democracy* in the context of CS and leverage our analysis of Justice Pursuit to show that random measurements are democratic. We conclude with a brief discussion of the broader role of democracy in CS.

In **Part IV** we turn to the problem of directly processing compressive measurements to filter or extract desired information. We begin in **Chapter 9** with an analysis of three fundamental signal processing problems: detection, classification,

and estimation. In the case of signal detection and classification from random measurements in the presence of Gaussian noise, we derive the optimal detector/classifier and analyze its performance. We show that in the high SNR regime we can reliably detect/classify with far fewer measurements than are required for recovery. We also propose a simple and efficient approach to the estimation of linear functions of the signal from random measurements. We argue that in all of these settings, we can exploit sparsity and random measurements to enable the design of efficient, universal acquisition hardware. While these choices do not exhaust the set of canonical signal processing operations, we believe that they provide a strong initial foundation for CSP.

Chapter 10 then analyzes the problem of filtering compressive measurements. We begin with a simple method for suppressing sparse interference. We demonstrate the relationship between this method and a key step in orthogonal greedy algorithms and illustrate its application to the problem of signal recovery in the presence of interference, or equivalently, signal recovery with partially known support. We then generalize this method to more general filtering methods, with a particular focus on the cancellation of bandlimited, but not necessarily sparse, interference.

We conclude with a summary of our findings, discussion of ongoing work, and directions for future research in **Chapter 11**.

This thesis is the culmination of a variety of intensive collaborations. Where appropriate, the first page of each chapter provides a footnote listing primary collaborators, who share credit for this work.

Part I

Sparse Signal Models

Chapter 2

Overview of Sparse Models

2.1 Mathematical Preliminaries

2.1.1 Vector spaces

Throughout this thesis, we will treat signals as real-valued functions having domains that are either continuous or discrete, and either infinite or finite. These assumptions will be made clear as necessary in each section. In the case of a discrete, finite domain, we can view our signals as vectors in N-dimensional Euclidean space, denoted by $\mathbb{R}^N$. We will denote the standard inner product in Euclidean space as

$$\langle x, y \rangle = y^T x = \sum_{i=1}^{N} x_i y_i.$$

We will also make frequent use of the ℓ_p norms, which are defined for $p \in [1, \infty]$ as

$$\|x\|_p = \begin{cases} \left(\sum_{i=1}^{N} |x_i|^p \right)^{\frac{1}{p}} & p \in [1, \infty); \\ \max_{i=1,2,\ldots,N} |x_i|, & p = \infty. \end{cases} \tag{2.1}$$

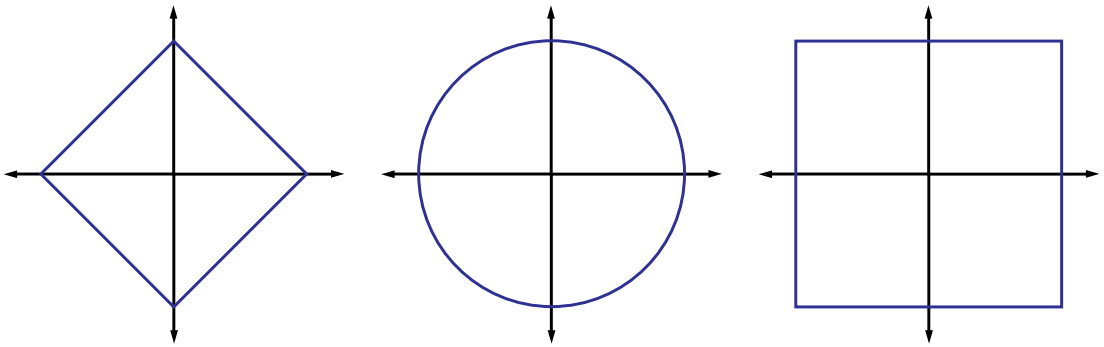

Figure 2.1: Unit balls in $\mathbb{R}^2$ for the ℓ_p norms for $p = 1, 2, \infty$

The ℓ_p norms have notably different properties for different values of p. To illustrate this, we show the unit ball, i.e., $\{x : \|x\|_p = 1\}$, induced by each of these norms in $\mathbb{R}^2$ in Figure 2.1.

We typically use norms as a measure of the strength of a signal, or the size of an error. For example, suppose we are given a signal $x \in \mathbb{R}^2$ and wish to approximate it using a point in a one-dimensional subspace A. If we measure the approximation error using an ℓ_p norm, then our task is to find the $\widehat{x} \in A$ that minimizes $\|x - \widehat{x}\|_p$. The choice of p will have a significant effect on the properties of the resulting approximation error. An example is illustrated in Figure 2.2. To compute the closest point in A to x using each ℓ_p norm, we can imagine growing an ℓ_p ball centered on x until it intersects with A. This will be the point $\widehat{x} \in A$ that is closest to x in the ℓ_p norm. We observe that larger p tends to spread out the error more evenly among the two coefficients, while smaller p leads to an error that is more unevenly distributed and tends to be sparse. This intuition generalizes to higher dimensions, and plays an important role in the development of the theory of CS.

Finally, in some contexts it is useful to extend the notion of ℓ_p norms to the case where $p < 1$. In this case, the "norm" defined in (2.1) fails to satisfy the triangle inequality, so it is actually a quasinorm. Moreover, we will also make frequent use of the notation $\|x\|_0 := |\text{supp}(x)|$, where $\text{supp}(x) = \{i : x_i \neq 0\}$ denotes the support of

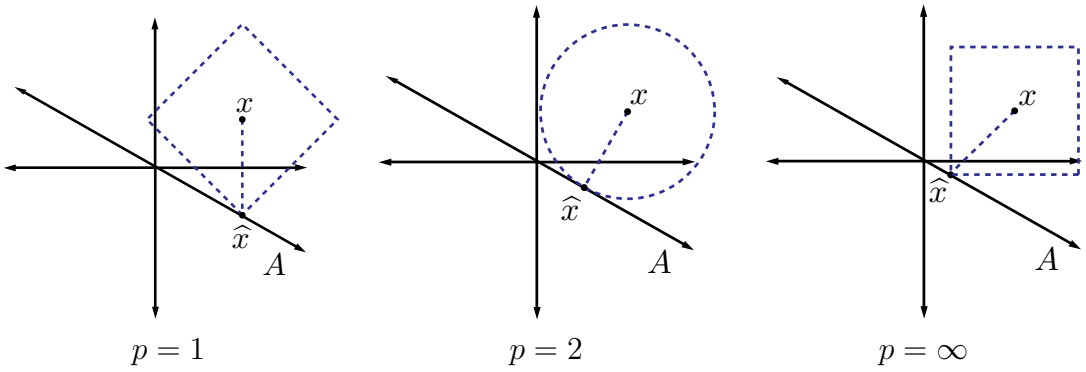

Figure 2.2: Best approximation of a point in $\mathbb{R}^2$ by a a one-dimensional subspace using the ℓ_p norms for $p = 1, 2, \infty$.

x. Note that $\|\cdot\|_0$ is not even a quasinorm, but one can easily show that

$$\lim_{p\to 0} \|x\|_p^p = |\mathrm{supp}(x)|,$$

justifying this choice of notation.

2.1.2 Bases

A set $\{\psi_i\}_{i=1}^N$ is called a basis for $\mathbb{R}^N$ if the vectors in the set span $\mathbb{R}^N$ and are linearly independent.[1] This implies each vector in the space has a unique representation as a linear combination of these basis vectors. Specifically, for any $x \in \mathbb{R}^N$, there exist (unique) coefficients $\{\alpha_i\}_{i=1}^N$ such that

$$x = \sum_{i=1}^{N} \alpha_i \psi_i.$$

Note that if we let Ψ denote the $N \times N$ matrix with columns given by ψ_i and let α denote the length-N vector with entries α_i, then we can represent this more compactly

[1] In any N-dimensional vector space, a basis will always consist of exactly N vectors, since fewer vectors are not sufficient to span the space, while if we add any additional vectors they are guaranteed to be linearly dependent on some subset of existing basis elements.

as

$$x = \Psi\alpha.$$

An important special case of a basis is an orthonormal basis (ONB), defined as a set of vectors $\{\psi_i\}_{i=1}^N$ that form a basis and whose elements are orthogonal and have unit norm, meaning that

$$\langle \psi_i, \psi_j \rangle = \begin{cases} 1, & i = j; \\ 0, & i \neq j. \end{cases}$$

An ONB has the advantage that the coefficients α can be easily calculated as

$$\alpha_i = \langle x, \psi_i \rangle,$$

or

$$\alpha = \Psi^T x$$

in matrix notation. This can easily be verified since the orthonormality of the columns of Ψ means that $\Psi^T\Psi = I$, where I denotes the $N \times N$ identity matrix.

2.1.3 Notation

Before proceeding, we will set the remainder of our notation. We will use $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ denote the ceiling and floor operators, respectively. We will use log throughout to denote the natural logarithm. When it is necessary to refer to logarithms of other bases, we will indicate this explicitly via a subscript as in $\log_{10}$. When taking a real number as an argument, $|x|$ denotes the absolute value of x, but when taking a set Λ as an argument, $|\Lambda|$ denotes the cardinality of Λ. By $x|_\Lambda$ we mean the length $|\Lambda|$ vector containing the entries of x indexed by Λ. When $\Lambda \subset \{1, 2, \ldots, N\}$ we let $\Lambda^c = \{1, 2, \ldots, N\} \backslash \Lambda$.

We will let $\mathcal{N}(\Phi)$ denote the null space of a matrix Φ, and $\mathcal{R}(\Phi)$ the range, or

column space, of Φ. By Φ_Λ we mean the $M \times |\Lambda|$ matrix obtained by selecting the columns of Φ indexed by Λ. We will assume throughout that when $|\Lambda| \leq M$, Φ_Λ is full rank, in which case we let $\Phi_\Lambda^\dagger := (\Phi_\Lambda^T \Phi_\Lambda)^{-1} \Phi_\Lambda^T$ denote the Moore-Penrose pseudoinverse of Φ_Λ.

We denote the orthogonal projection operator onto $\mathcal{R}(\Phi)$ by $P_\Phi = \Phi\Phi^\dagger$. When considering projections onto $\mathcal{R}(\Phi_\Lambda)$, we will also use the simpler notation P_Λ in place of P_{Φ_Λ}. Similarly, $P_\Lambda^\perp = (I - P_\Lambda)$ is the orthogonal projection operator onto the orthogonal complement of $\mathcal{R}(\Phi_\Lambda)$. We note that any orthogonal projection operator P obeys $P = P^T = P^2$.

Finally, we will let $\mathbb{P}(\text{event})$ denote the probability of a given event, and we will let

$$\mathbb{E}(g(X)) = \int_{-\infty}^{\infty} g(x) f(x)\, dx$$

denote the expected value of $g(X)$, where X is a random variable with probability density function $f(x)$ defined on $\mathbb{R}$.

2.2 Sparse Signals

2.2.1 Sparsity and nonlinear approximation

Sparse signal models provide a mathematical framework for capturing the fact that in many cases these high-dimensional signals contain relatively little information compared to their ambient dimensionality. Sparsity has long been exploited in signal processing and approximation theory for tasks such as compression [12] and denoising [8], and in statistics and learning theory as a method for avoiding overfitting [64]. Sparsity can be thought of as one incarnation of *Occam's razor* — when faced with many possible ways to represent a signal, the simplest choice is the best one.

Mathematically, we say that a signal x is K-sparse when it has at most K nonzeros,

i.e., $\|x\|_0 \leq K$. We let

$$\Sigma_K = \{x : \|x\|_0 \leq K\} \tag{2.2}$$

denote the set of all K-sparse signals. Typically, we will be dealing with signals that are not themselves sparse, but which admit a sparse representation in some basis Ψ. In this case we will still refer to x as being K-sparse, with the understanding that we can express x as $x = \Psi\alpha$ where $\|\alpha\|_0 \leq K$. When necessary, we will use $\Psi(\Sigma_K)$ to denote the set of all signals that are K-sparse when represented in the basis Ψ.

As a traditional application of sparse models, we consider the problems of image compression and image denoising. Most natural images are characterized by large smooth or textured regions and relatively few sharp edges. Signals with this structure are known to be very nearly sparse when represented using a multiscale wavelet transform [9]. The wavelet transform consists of recursively dividing the image into its low- and high-frequency components. The lowest frequency components provide a coarse scale approximation of the image, while the higher frequency components fill in the detail and resolve edges. Wavelet coefficients can be grouped into a tree-like structure as shown in Figure 1.1. What we see when we compute a wavelet transform of a typical natural image is that most coefficients are very small. Hence, we can obtain a good approximation to the signal by setting the small coefficients to zero, or *thresholding* the coefficients, to obtain a sparse representation of the image. Figure 2.3 shows an example of such an image and its K-term approximation. This is the heart of nonlinear approximation [12] — nonlinear because the choice of which coefficients to keep in the approximation depends on the signal itself. Similarly, given the knowledge that natural images are approximately sparse, this same thresholding operation serves as an effective method for rejecting certain kinds of signal noise [8].

(a) (b)

Figure 2.3: Sparse approximation of a natural image. (a) Original image (b) Approximation to image obtained by keeping only the largest 10% of the wavelet coefficients.

2.2.2 Geometry of sparse signals

Sparsity is a highly nonlinear signal model. This can be seen by observing that given a pair of K-sparse signals, a linear combination of the two signals will in general no longer be K sparse, since their supports may not overlap. That is, for any $x, y \in \Sigma_K$, we do not necessarily have that $x + y \in \Sigma_K$ (although we do have that $x + y \in \Sigma_{2K}$). This is illustrated in Figure 2.4, which shows Σ_2 embedded in $\mathbb{R}^3$, i.e., the set of all 2-sparse signals in $\mathbb{R}^3$.

While the set of sparse signals Σ_K does not form a linear space, it does satisfy a great deal of structure. Specifically, it consists of the union of all possible $\binom{N}{K}$ subspaces. In Figure 2.4 we have only $\binom{3}{2} = 3$ possible subspaces, but for larger values of N and K we must consider a potentially huge number of subspaces. This will have significant algorithmic consequences in the development of the algorithms for sparse approximation and sparse recovery described below.

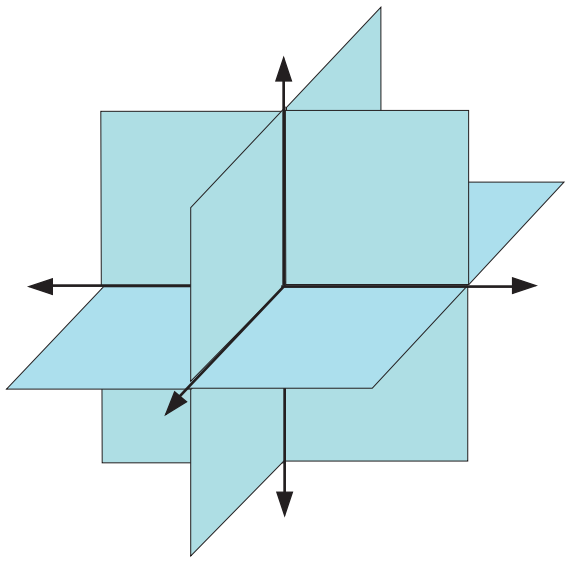

Figure 2.4: Union of subspaces defined by $\Sigma_2 \subset \mathbb{R}^3$, i.e., the set of all 2-sparse signals in $\mathbb{R}^3$.

2.2.3 Compressible signals

An important point in practice is that few real-world signals are *truly* sparse; rather they are compressible, meaning that they can be well-approximated by a sparse signal. We can quantify this by calculating the error incurred by approximating a signal x by some $\widehat{x} \in \Sigma_K$:

$$\sigma_K(x)_p = \min_{\widehat{x} \in \Sigma_K} \|x - \widehat{x}\|_p. \tag{2.3}$$

If $x \in \Sigma_K$, then clearly $\sigma_K(x)_p = 0$ for any p. Moreover, one can easily show that the thresholding strategy described above (keeping only the K largest coefficients) results in the optimal approximation as measured by (2.3) for all ℓ_p norms.

Another way to think about compressible signals is to consider the rate of decay of their coefficients. For many important classes of signals there exist bases such that the coefficients obey a power law decay, in which case the signals are highly compressible. Specifically, if $x = \Psi\alpha$ and we sort the coefficients α_i such that $|\alpha_1| \geq |\alpha_2| \geq \cdots \geq |\alpha_N|$, then we say that the coefficients obey a power law decay if there

exist constants $C_1, q > 0$ such that

$$|\alpha_i| \leq C_1 i^{-q}.$$

The larger q is, the faster the magnitudes decay, and the more compressible a signal is. Because the magnitudes of their coefficients decay so rapidly, compressible signals can be represented accurately by $K \ll N$ coefficients. Specifically, for such signals there exist constants $C_2, r > 0$ depending only on C_1 and q such that

$$\sigma_K(x)_2 \leq C_2 K^{-r}.$$

In fact, one can show that $\sigma_K(x)_2$ will decay as K^{-r} if and only if the sorted coefficients α_i decay as $i^{-r+1/2}$ [12].

2.3 Compressive Sensing

CS has emerged as a new framework for signal acquisition and sensor design that enables a potentially large reduction in the cost of acquiring signals that have a sparse or compressible representation [13–19]. Specifically, given a signal $x \in \mathbb{R}^N$, we consider measurement systems that acquire M linear measurements.[2] We can represent this process mathematically as

$$y = \Phi x, \tag{2.4}$$

[2]Note that the standard CS framework assumes that x is a finite-length, discrete-time vector, while in practice we will be interested in designing measurement systems for acquiring analog, continuous-time signals. We will discuss how to extend this discrete formulation to the continuous-time case in greater detail in Chapter 5, but for now we will just think of x as a finite-length window of Nyquist-rate samples, and we will see later how to directly acquire compressive measurements without first sampling at the Nyquist-rate.

where Φ is an $M \times N$ matrix and $y \in \mathbb{R}^M$. The matrix Φ represents a *dimensionality reduction*, i.e., it maps $\mathbb{R}^N$, where N is generally large, into $\mathbb{R}^M$, where M is typically much smaller than N. In this case, we refer to the measurements y as *compressive measurements*.

There are two main theoretical questions in CS. First, how should we design Φ to ensure that it preserves the information in the signal x? Secondly, how can we recover the original signal x from the measurements y? In the absence of some additional information concerning x, the answer is straightforward: we must ensure that Φ is invertible, in which case we can simply recover the original signal via $x = \Phi^{-1}y$. Unfortunately, this requires full measurements (setting $M = N$). In the case where our data is sparse or compressible, the answers change dramatically. We will be able to design matrices Φ with $M \ll N$ and be able to recover the original signal accurately and efficiently using a variety of practical algorithms.

We will address the question of how to design Φ in detail in Part II, but before we do so it will be useful to first gain some intuition into how we will solve the second problem of signal recovery. The challenge here is to somehow exploit the fact that x lives in or near Σ_K, or $\Psi(\Sigma_K)$ for some known basis Ψ. In the former case, we have that y is a linear combination of at most K columns of Φ, while in the latter case y is a combination of at most K columns of the matrix $\widetilde{\Phi} = \Phi\Psi$. Without loss of generality, we will restrict our attention to the case where $x \in \Sigma_K$, since we can always reduce the problem to this case via a simple substitution. Fortunately, the sparse recovery problem has received significant attention over the years in the context of computing sparse representations with overcomplete dictionaries.

2.4 Computing Optimal Sparse Representations

Suppose that we are given a vector y and wish to represent y as a sparse linear combination of the columns of an $M \times N$ matrix Φ. If $M = N$, then the answer is trivial — Φ represents a basis, and so we simply compute the expansion coefficients α and then threshold them to obtain the optimal sparse representation as described in Section 2.2.3. The challenge arises when $M < N$. In this case, Φ is no longer a basis, but rather an overcomplete dictionary, with the consequence that there is (in general) no unique set of expansion coefficients. If we want to find the optimal sparse representation, then we must somehow find the most compressible expansion, or equivalently, we must search through all possible sets of K columns to find the best K-sparse representation. Unfortunately, there are $\binom{N}{K}$ possibilities, and so this strategy becomes impossible for even extremely modest values of K and N.

In response to this challenge, over the years there have been various algorithms and heuristics that have been proposed for solving this and closely related problems in signal processing, statistics, and computer science. We now provide a brief overview of some of the key algorithms that we will make use of in this thesis. We refer the reader to [46] and references therein for a more thorough survey of these methods.

2.4.1 Basis Pursuit and optimization-based methods

We can formulate the sparse approximation problem as a nonconvex optimization problem. Specifically, we would like to solve the problem

$$\widehat{x} = \arg\min_{x} \ \|x\|_0 \quad \text{subject to} \quad \Phi x = y, \tag{2.5}$$

i.e., we would like to find the sparsest x such that $y = \Phi x$. Of course, even if x is truly K-sparse, adding even a small amount of noise to y will result in a solution to (2.5) with M nonzeros, rather than K. To introduce some tolerance for noise and other

errors, as well as robustness to compressible rather than sparse signals, we would typically rather solve a slight variation to (2.5):

$$\widehat{x} = \arg\min_x \ \|x\|_0 \quad \text{subject to} \quad \|\Phi x - y\|_2 \leq \epsilon. \tag{2.6}$$

Note that we can express (2.6) in two alternative, but equivalent, manners. While in many cases the choice for the parameter ϵ may by clear, in other cases it may be more natural to specify a desired level of sparsity K. In this case we can consider the related problem of

$$\widehat{x} = \arg\min_x \ \|\Phi x - y\|_2 \quad \text{subject to} \quad \|x\|_0 \leq K. \tag{2.7}$$

Alternatively, we can also consider the unconstrained version of this problem:

$$\widehat{x} = \arg\min_x \ \|x\|_0 + \lambda \|\Phi x - y\|_2. \tag{2.8}$$

Of course, all of these formulations are nonconvex, and hence potentially very difficult to solve. In fact, one can show that for a general dictionary Φ, even finding a solution that approximates the true minimum is NP-hard [62].

The difficulty in solving these problems arises from the fact that $\|\cdot\|_0$ is a nonconvex function. Thus, one avenue for translating these problems into something more tractable is to replace $\|\cdot\|_0$ with its convex relaxation $\|\cdot\|_1$. Thus, in the case of (2.5) this yields

$$\widehat{x} = \arg\min_x \ \|x\|_1 \quad \text{subject to} \quad y = \Phi x, \tag{2.9}$$

and in the case of (2.6) we obtain

$$\widehat{x} = \arg\min_x \ \|x\|_1 \quad \text{subject to} \quad \|\Phi x - y\|_2 \leq \epsilon. \tag{2.10}$$

We will refer to (2.9) as *Basis Pursuit* (BP), and (2.10) as *Basis Pursuit De-Noising* (BPDN), following the terminology introduced in [65]. These problems can be posed as linear programs and solved using a variety of methods.

Before discussing an example of one such method, we note that the use of ℓ_1 minimization to promote or exploit sparsity has a long history, dating back at least to the work of Beurling on Fourier transform extrapolation from partial observations [29]. In a somewhat different context, in 1965 Logan [66] showed that a bandlimited signal can be perfectly recovered in the presence of *arbitrary* corruptions on a small interval. Again, the recovery method consists of searching for the bandlimited signal that is closest to the observed signal in the ℓ_1 norm. This can be viewed as further validation of the intuition gained from Figure 2.2 — the ℓ_1 norm is well-suited to sparse errors.

The use of ℓ_1 minimization on large problems finally became practical with the explosion of computing power in the late 1970's and early 1980's. In one of its first practical applications, it was demonstrated that geophysical signals consisting of spike trains could be recovered from only the high-frequency components of these signals by exploiting ℓ_1 minimization [67–69]. Finally, in the 1990's there was a renewed interest in these approaches within the signal processing community for the purpose of finding sparse approximations to signals and images when represented in overcomplete dictionaries or unions of bases [9, 65]. Meanwhile, the ℓ_1 variant of (2.7) began to receive significant in the statistics literature as a method for variable selection in regression known as LASSO [70].

Finally, we conclude with an illustrative example of an algorithm known as Fixed-Point Continuation (FPC) which is designed to solve the ℓ_1 variant of (2.8) [40]. This approach is an iterative, gradient descent method that will bear a great deal of similarity to some of the greedy methods described below, but which can actually be proven to converge to the ℓ_1 optimum. The heart of the algorithm is a gradient

descent step on the quadratic penalty term. Specifically, note that

$$\nabla \|y - \Phi x\|_2^2 = 2\Phi^T(y - \Phi x).$$

At the ℓ^{th} iteration of the algorithm, we have an estimate x^ℓ, and thus the gradient descent step would consist of

$$x^{\ell+1} = x^\ell - \tau \Phi^T(y - \Phi x^\ell),$$

where τ is a parameter the user must set specifying the step size. This gradient descent step is then followed by *soft thresholding* to complete the iteration. The full algorithm for FPC is specified in Algorithm 1. We use $\text{soft}(x, \alpha)$ to denote the soft thresholding, or shrinkage, operator:

$$[\text{soft}(x, \alpha)]_i = \begin{cases} x_i - \alpha, & x_i > \alpha; \\ 0, & x_i \in [-\alpha, \alpha]; \\ x_i + \alpha, & x_i < \alpha. \end{cases} \tag{2.11}$$

In our statement of the algorithm, we use r^ℓ to denote the *residual* $y - \Phi x^\ell$ and refer to the step of computing $h^\ell = \Phi^T r^\ell$ as the step of computing the *proxy*, for reasons that will become clear when we draw parallels between FPC and the greedy algorithms described below.

2.4.2 Greedy algorithms and iterative methods

While convex optimization techniques like FPC are powerful methods for computing sparse representations, there are also a variety of greedy/iterative methods for solving such problems. We emphasize that in practice, many ℓ_1 solvers are themselves iterative algorithms, and in fact we will see that FPC is remarkably similar to some of

Algorithm 1 Fixed-Point Continuation (FPC)

input: Φ, y, λ, τ, stopping criterion
initialize: $r^0 = y$, $x^0 = 0$, $\ell = 0$
while not converged **do**
  proxy: $h^\ell = \Phi^T r^\ell$
  update: $x^{\ell+1} = \text{soft}(x^\ell - \tau h^\ell, \tau/\lambda)$
  $r^{\ell+1} = y - \Phi x^{\ell+1}$
  $\ell = \ell + 1$
end while
output: $\widehat{x} = x^\ell$

the algorithms discussed below. However, the fundamental difference is that FPC is actually proven to optimize an ℓ_1 objective function, while the methods below mostly arose historically as simple heuristics that worked well in practice and do not claim to optimize any such objective function. We will see later, however, that many of these algorithms actually have similar performance guarantees to those of the seemingly more powerful optimization-based approaches.

We begin with the oldest of these algorithms. *Matching Pursuit* (MP), shown in Algorithm 2, provides the basic structure for all of the greedy algorithms to follow [71]. In the signal processing community, MP dates back to [71], but essentially the same algorithm had been independently developed in a number of other fields even earlier. In the context of feature selection for linear regression, the algorithm of *Forward Selection* is nearly identical to MP [72, 73], as well as the onion peeling algorithms for multiuser detection in digital communications [74].

At the beginning of each iteration of MP, $r^\ell = y - \Phi x^\ell$ represents the residual, or the part of y that we have not yet explained using our estimate of x. Each iteration then forms the estimate $h^\ell = \Phi^T r^\ell$, which serves as a proxy, or very rough estimate, of the part of x we have yet to identify. At this point, each algorithm will perform an update using this proxy vector. A common theme among greedy algorithms is the use of *hard thresholding* (as opposed to soft thresholding, which commonly arises in the optimization-based approaches) to keep track of an index or indices that we wish

Algorithm 2 Matching Pursuit (MP)

input: Φ, y, stopping criterion
initialize: $r^0 = y$, $x^0 = 0$, $\ell = 0$
while not converged **do**
  proxy: $h^\ell = \Phi^T r^\ell$
  update: $x^{\ell+1} = x^\ell + \text{hard}(h^\ell, 1)$
  $r^{\ell+1} = y - \Phi x^{\ell+1}$
  $\ell = \ell + 1$
end while
output: $\widehat{x} = x^\ell$

to update. Specifically, we will define[3]

$$[\text{hard}(x, K)]_i = \begin{cases} x_i, & |x_i| \text{ is among the } K \text{ largest elements of } |x|; \\ 0, & \text{otherwise.} \end{cases} \tag{2.12}$$

MP applies hard directly to the proxy vector h^ℓ to pick a single coefficient to update, and then simply uses the value of h^ℓ on that coefficient as the update step.

MP is also closely related to the more recently developed algorithm of *Iterative Hard Thresholding* (IHT) [30]. The only difference between MP and IHT is that we replace the update step $x^{\ell+1} = x^\ell + \text{hard}(h^\ell, 1)$ with

$$x^{\ell+1} = \text{hard}\left(x^\ell + h^\ell, K\right). \tag{2.13}$$

This change allows IHT to be more aggressive at the beginning of the algorithm, but ensures that after many iterations, $\|x^\ell\|_0$ remains well-controlled.

However, the most common extension of MP is *Orthogonal Matching Pursuit* (OMP) [71, 75, 76]. The algorithm, provided in Algorithm 3, is only slightly different than MP. Both algorithms begin by forming the proxy vector h^ℓ and then identify-

[3]Note that we have defined our thresholding operator not by assigning it a threshold value, as we did in (2.11), but by dictating the number of elements we wish to keep. This is to ensure that $|\text{supp}(\text{hard}(x, K))| = K$. In the event that there are ties among the $|x_i|$, we do not specify which x_i are kept. The algorithm designer is free to choose any tiebreaking method available.

Algorithm 3 Orthogonal Matching Pursuit (OMP)

input: Φ, y, stopping criterion
initialize: $r^0 = y$, $x^0 = 0$, $\Lambda^0 = \emptyset$, $\ell = 0$
while not converged **do**
  proxy: $h^\ell = \Phi^T r^\ell$
  identify: $\Lambda^{\ell+1} = \Lambda^\ell \cup \operatorname{supp}\left(\operatorname{hard}(h^\ell, 1)\right)$
  update: $x^{\ell+1} = \arg\min_{z:\ \operatorname{supp}(z) \subseteq \Lambda^{\ell+1}} \|y - \Phi z\|_2$
  $r^{\ell+1} = y - \Phi x^{\ell+1}$
  $\ell = \ell + 1$
end while
output: $\widehat{x} = x^\ell$

ing the largest component of h^ℓ. However, where MP simply uses the thresholded version of h^ℓ as the signal update, OMP does something more sophisticated — it finds the least-squares optimal recovery among all signals living on the support of the coefficients chosen in the first ℓ iterations. One can show that this will ensure that once a particular coefficient has been selected, it will never be selected again in a later iteration. Thus, we always have that $\|x^\ell\|_0 = \ell$. Moreover, the output $\widehat{x}$ will be the least-squares optimal recovery among all signals living on $\operatorname{supp}(\widehat{x})$. These properties come at an increased computational cost per iteration over MP, but in practice this additional computational cost can be justified, especially if it results in a more accurate recovery and/or fewer iterations.

2.4.3 Picky Pursuits

In recent years, there have been a great many variants of OMP which have been proposed and studied [33, 35, 37, 42–44]. These algorithms share many of the same features. First, they all modify the identification step by selecting more than one index to add to the active set Λ^ℓ at each iteration. In the case of *Stagewise Orthogonal Matching Pursuit* (StOMP) [37], this is the only difference. The different approaches vary in this step — some choose all of the coefficients that exceed some pre-specified threshold, while others pick cK at a time for some parameter c.

In a sense, these algorithms seem *more* greedy than OMP, since they potentially add more than just one coefficient to Λ^ℓ at a time. However, as these algorithms were developed they began to incorporate another feature — the ability to remove coefficients from Λ^ℓ. While this capability was not present in StOMP, it began to appear in *Regularized Orthogonal Matching Pursuit* (ROMP) [42, 43], which followed StOMP in adding multiple indices at once, but carefully selected these indices to ensure that they were comparable in size. *Compressive Sampling Matching Pursuit* (CoSaMP) [44], *Subspace Pursuit* (SP) [35], and *DThresh* [33] take this one step further by explicitly removing indices from Λ^ℓ at each iteration. While these algorithms are still typically referred to as greedy algorithms, they are actually quite picky in which coefficients they will retain after each iteration.

As an illustration of these algorithms, we will describe ROMP and CoSaMP in some detail. We first briefly describe the difference between ROMP and OMP, which lies only in the identification step: whereas OMP adds only one index to Λ^ℓ at each iteration, ROMP adds up to K indices to Λ^ℓ at each iteration. Specifically, ROMP first selects the indices corresponding to the K largest elements in magnitude of h^ℓ (or all nonzero elements of h^ℓ if h^ℓ has fewer than K nonzeros), and denotes this set as Ω^ℓ. The next step is to *regularize* this set so that the values are comparable in magnitude. To do this, we define $R(\Omega^\ell) := \{\Omega \subseteq \Omega^\ell : |h_i^\ell| \leq 2|h_j^\ell| \ \forall i, j \in \Omega\}$, and set

$$\Omega_0^\ell := \underset{\Omega \in R(\Omega^\ell)}{\arg\max} \|h^\ell|_\Omega\|_2,$$

i.e., Ω_0^ℓ is the set with maximal energy among all regularized subsets of Ω^ℓ. Setting $\Lambda^{\ell+1} = \Lambda^\ell \cup \Omega_0^\ell$, the remainder of the ROMP algorithm is identical to OMP. The full algorithm is shown in Algorithm 4.

Finally, we conclude with a discussion of CoSaMP. CoSaMP, shown in Algorithm 5, differs from OMP both in the identification step *and* in the update step. At each

Algorithm 4 Regularized Orthogonal Matching Pursuit (ROMP)

input: Φ, y, K, stopping criterion
initialize: $r^0 = y$, $x^0 = 0$, $\Lambda^0 = \emptyset$, $\ell = 0$
while not converged **do**
  proxy: $h^\ell = \Phi^T r^\ell$
  identify: $\Omega^\ell = \text{supp}(\text{hard}(h^\ell, K))$
  $\Lambda^{\ell+1} = \Lambda^\ell \cup \text{regularize}(\Omega^\ell)$
  update: $x^{\ell+1} = \arg\min_{z:\ \text{supp}(z)\subseteq\Lambda^{\ell+1}} \|y - \Phi z\|_2$
  $r^{\ell+1} = y - \Phi x^{\ell+1}$
  $\ell = \ell + 1$
end while
output: $\widehat{x} = x^\ell = \arg\min_{z:\ \text{supp}(z)\subseteq\Lambda^\ell} \|y - \Phi z\|_2$

Algorithm 5 Compressive Sampling Matching Pursuit (CoSaMP)

input: Φ, y, K, stopping criterion
initialize: $r^0 = y$, $x^0 = 0$, $\Lambda^0 = \emptyset$, $\ell = 0$
while not converged **do**
  proxy: $h^\ell = \Phi^T r^\ell$
  identify: $\Lambda^{\ell+1} = \text{supp}(x^\ell) \cup \text{supp}(\text{hard}(h^\ell, 2K))$
  update: $\widetilde{x} = \arg\min_{z:\ \text{supp}(z)\subseteq\Lambda^{\ell+1}} \|y - \Phi z\|_2$
  $x^{\ell+1} = \text{hard}(\widetilde{x}, K)$
  $r^{\ell+1} = y - \Phi x^{\ell+1}$
  $\ell = \ell + 1$
end while
output: $\widehat{x} = x^\ell = \arg\min_{z:\ \text{supp}(z)\subseteq\Lambda^\ell} \|y - \Phi z\|_2$

iteration the algorithm begins with an x^ℓ with at most K nonzeros. It then adds $2K$ indices to Λ^ℓ.[4] At this point, $|\Lambda^\ell| \leq 3K$. Proceeding as in OMP, CoSaMP solves a least-squares problem to update x^ℓ, but rather than updating with the full solution to the least-squares problem, which will have up to $3K$ nonzeros, CoSaMP thresholds this solution and updates with a pruned version so that x^ℓ will have only K nonzeros.

[4]The choice of $2K$ is primarily driven by the proof technique, and is not intended to be interpreted as an optimal or necessary choice. For example, in [35] it is shown that the choice of K is sufficient to establish similar performance guarantees to those for CoSaMP.

Part II

Sparse Signal Acquisition

Chapter 3

Stable Embeddings of Sparse Signals

One of the core problems in signal processing concerns the acquisition of a discrete, digital representation of a signal. Mathematically, we can represent an acquisition system that obtains M linear measurements as an operator $\Phi : \mathcal{X} \to \mathbb{R}^M$, where $\mathcal{X}$ represents our signal space. For example, in classical sampling systems we assume that $\mathcal{X}$ is the set of all bandlimited signals, in which case the Nyquist-Shannon sampling theorem dictates that acquiring uniform samples in time at the Nyquist rate is optimal, in the sense that it exactly preserves the information in the signal, and that with any fewer samples there would be some signals in our model which we would be unable to recover.

In CS, we extend our concept of a measurement system to allow general linear operators Φ, not just sampling systems. As in the classical setting, we wish to design our measurement system Φ with two competing goals: *(i)* we want to acquire as few measurements as possible, i.e., we want M to be small, and *(ii)* we want to ensure that we preserve the information in our signal. While there are many possible ways to mathematically capture the notion of information preservation, a simple yet

powerful approach is to require that Φ be a *stable embedding* of the signal model. Specifically, given a distance metric $d_{\mathcal{X}}(x, y)$ defined on pairs $x, y \in \mathcal{X}$, then Φ is a stable embedding of $\mathcal{X}$ if there exists a constant $\delta \in (0, 1)$ such that

$$(1 - \delta) d_{\mathcal{X}}(x, y) \leq \|\Phi x - \Phi y\|_{\ell_p} \leq (1 + \delta) d_{\mathcal{X}}(x, y) \tag{3.1}$$

for all $x, y \in \mathcal{X}$. An operator satisfying (3.1) is also called *bi-Lipschitz*. This property ensures that signals that are well-separated in $\mathcal{X}$ remain well-separated after the application of Φ. This implies that Φ is one-to-one, and hence invertible — moreover, in the case where the measurements are perturbed, this also guarantees a degree of stability.

In this chapter[1] we focus on the special case where $\mathcal{X} = \Sigma_K \subset \mathbb{R}^N$ and we measure distances with the ℓ_2 norm, in which case the property in (3.1) is also known as the *restricted isometry property* (RIP) or *uniform uncertainty principle* (UUP). We examine the role that the RIP plays in the stability of sparse recovery algorithms, showing that in certain settings it is actually a necessary condition. We then provide a brief overview of the theoretical implications of the RIP for some of the sparse recovery algorithms described in Section 2.4. We then close by establishing a pair of lower bounds on the number of measurements M that any matrix satisfying the RIP must satisfy.

3.1 The Restricted Isometry Property

In [77], Candès and Tao introduced the following isometry condition on matrices Φ and established its important role in CS.

[1]Thanks to Peter G. Binev and Piotr Indyk for many useful discussions and helpful suggestions, especially regarding Theorem 3.5.

Definition 3.1. We say that a matrix Φ satisfies the *restricted isometry property* (RIP) of order K if there exists a $\delta_K \in (0,1)$ such that

$$(1-\delta_K)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1+\delta_K)\|x\|_2^2, \tag{3.2}$$

for all $x \in \Sigma_K$.

While the inequality in (3.2) may appear to be somewhat different from our notion of a stable embedding in (3.1), they are in fact equivalent. Specifically, if we set $\mathcal{X} = \Sigma_K$ and $d_{\mathcal{X}}(x,y) = \|x-y\|_2$, then Φ is a stable embedding of Σ_K if and only if Φ satisfies the RIP of order $2K$ (since for any $x, y \in \Sigma_K$, $x - y \in \Sigma_{2K}$).

Note that if Φ satisfies the RIP of order K with constant δ_K, then for any $K' < K$ we automatically have that Φ satisfies the RIP of order K' with constant $\delta_{K'} \leq \delta_K$. Moreover, in [44] it is shown (Corollary 3.4) that if Φ satisfies the RIP of order K with a sufficiently small constant, then it will also automatically satisfy the RIP of order cK for certain c, albeit with a somewhat worse constant.

Lemma 3.1 (Needell-Tropp [44])**.** *Suppose that Φ satisfies the RIP of order K with constant δ_K. Let c be a positive integer. Then Φ satisfies the RIP of order $K' = c\left\lfloor \frac{K}{2} \right\rfloor$ with constant $\delta_{K'} < c\,\delta_K$.*

This lemma is trivial for $c = 1, 2$, but for $c \geq 3$ (and $K \geq 4$) this allows us to extend from RIP of order K to higher orders. Note however that δ_K must be sufficiently small in order for the resulting bound to be useful. In particular, this lemma only yields $\delta_{K'} < 1$ when $\delta_K < 1/c$. Thus, we cannot extend the RIP to arbitrarily large order. We will make use of this fact below in providing a lower bound on the number of measurements necessary to obtain a matrix satisfying the RIP with a particular constant. Note that when K is clear from the context, we will often omit the dependence of δ_K on K, and simply use δ to denote the RIP constant.

3.2 The RIP and Stability

We will see later that if a matrix Φ satisfies the RIP, then this is sufficient for a variety of algorithms to be able to successfully recover a sparse signal x from the measurements Φx. First, however, we will take a closer look at whether the RIP is actually necessary. It is clear that the lower bound in the RIP is a necessary condition if we wish to be able to recover all sparse signals x from the measurements Φx. Specifically, if x has at most K nonzero entries, then Φ must satisfy the lower bound of the RIP of order $2K$ with $\delta_{2K} < 1$ in order to ensure that any algorithm can recover x from the measurements y. To see this, observe that if Φ fails to satisfy the RIP for any $\delta_{2K} < 1$, then there exists a vector h such that h has at most $2K$ nonzeros and $\Phi h = 0$. Since h has at most $2K$ nonzeros, we can write $h = x + x'$, where both x and x' have at most K nonzeros. This yields $\Phi x = \Phi x'$, hence no method can ever hope to successfully recover *all* K-sparse signals.

Moreover, we can say even more about the necessity of the RIP by considering the following notion of stability.

Definition 3.2. Let $\Phi : \mathbb{R}^N \to \mathbb{R}^M$ denote a measurement matrix and $\Delta : \mathbb{R}^M \to \mathbb{R}^N$ denote a recovery algorithm. We say that the pair (Φ, Δ) is *C-stable* if for any $x \in \Sigma_K$ and any $e \in \mathbb{R}^M$ we have that

$$\|\Delta(\Phi x + e) - x\|_2 \leq C\|e\|_2.$$

This definition simply says that if we add a small amount of noise to the measurements, the impact of this on the recovered signal should not be arbitrarily large. The theorem below demonstrates that the existence of any decoding algorithm (however impractical) that can stably recover from noisy measurements necessitates that Φ satisfy the lower bound of (3.2) with a constant determined by C.

Theorem 3.1. *If the pair* (Φ, Δ) *is* C*-stable, then*

$$\frac{1}{C}\|x\|_2 \leq \|\Phi x\|_2 \tag{3.3}$$

for all $x \in \Sigma_{2K}$.

Proof. Pick any $x, y \in \Sigma_K$. Define

$$e_x = \frac{\Phi(y-x)}{2} \qquad \text{and} \qquad e_y = \frac{\Phi(x-y)}{2},$$

and note that

$$\Phi x + e_x = \Phi y + e_y = \frac{\Phi(x+y)}{2}.$$

Let $\widehat{x} = \Delta(\Phi x + e_x) = \Delta(\Phi y + e_y)$. From the triangle inequality and the definition of C-stability, we have that

$$\begin{aligned}
\|x-y\|_2 &= \|x - \widehat{x} + \widehat{x} - y\|_2 \\
&\leq \|x - \widehat{x}\|_2 + \|\widehat{x} - y\|_2 \\
&\leq C\|e_x\|_2 + C\|e_y\|_2 = 2C\left\|\frac{\Phi(x-y)}{2}\right\|_2 \\
&= C\|\Phi x - \Phi y\|_2.
\end{aligned}$$

Since this holds for any $x, y \in \Sigma_K$, the result follows. □

Note that as $C \to 1$, we have that Φ must satisfy the lower bound of (3.2) with $\delta_K = 1 - 1/C^2 \to 0$. Thus, if we desire to reduce the impact of noise in our recovered signal we must adjust Φ so that it satisfies the lower bound of (3.2) with a tighter constant.

One might respond to this result by arguing that since the upper bound is not necessary, we can avoid redesigning Φ simply by rescaling Φ so that as long as Φ satisfies the RIP with $\delta_{2K} < 1$, the rescaled version $A\Phi$ will satisfy (3.3) for any

constant C. In settings where the size of the noise is independent of our choice of Φ, this is a valid point — by scaling Φ we are essentially adjusting the gain on the "signal" part of our measurements, and if increasing this gain does not impact the noise, then we can achieve arbitrarily high signal-to-noise ratios, so that eventually the noise is negligible compared to the signal. However, in most practical settings the noise is not independent of Φ. For example, consider the case where the noise vector e represents quantization noise produced by a finite dynamic range quantizer with B bits. Suppose the measurements lie in the interval $[-T, T]$, and we have adjusted the quantizer to capture this range. If we rescale Φ by A, then the measurements now lie between $[-AT, AT]$ and we must scale the dynamic range of our quantizer by A. In this case the resulting quantization error is simply Ae, and we have achieved *no reduction* in the reconstruction error.

3.3 Consequences of the RIP

3.3.1 ℓ_1-norm minimization

We have argued that the RIP is in a certain sense a necessary condition for stability. In fact, we can show that for many of the sparse recovery algorithms described in Section 2.4, the RIP also provides a sufficient condition to guarantee robustness to both noise and to compressible signals. While there have been numerous variations in the analysis of ℓ_1 minimization applied to the problem of sparse recovery — in particular the BPDN formulation in (2.10) — we will choose a single result representative of the literature. See [78] for a short proof of this theorem.

Theorem 3.2 (Candès [78]). *Suppose that Φ satisfies the RIP of order $2K$ with isometry constant $\delta_{2K} < \sqrt{2}-1$. Given measurements of the form $y = \Phi x + e$, where $\|e\|_2 \leq \epsilon$, the solution to (2.10) obeys*

$$\|\widehat{x} - x\|_2 \leq C_0 \epsilon + C_1 \frac{\sigma_K(x)_1}{\sqrt{K}}, \tag{3.4}$$

where $\sigma_K(x)_1$ is defined as in (2.3) and where

$$C_0 = 4\frac{\sqrt{1+\delta_{2K}}}{1-(1+\sqrt{2})\delta_{2K}}, \quad C_1 = 2\frac{1-(1-\sqrt{2})\delta_{2K}}{1-(1+\sqrt{2})\delta_{2K}}.$$

Some comments are in order. First, note that the reconstruction error is bounded by two terms. The first term is determined by the bound on the measurement noise e. This tells us that if we add a small amount of noise to the measurements, its impact on the recovered signal remains well-controlled. Moreover, as the noise bound approaches zero, we see that the impact of the noise on the reconstruction error will also approach zero. The second term measures the error incurred by approximating the signal x as a K-sparse signal (where the error is measured using the ℓ_1 norm). In the event that x is compressible, then the error again remains well-controlled. Note that this term vanishes if x is perfectly K-sparse.[2] Moreover, in the event that x is K-sparse and there is no noise, then we obtain an *exact* recovery.

There have been many efforts to improve on the constants in Theorem 3.2 and to weaken the assumption on the constant δ, but most of this work results in theorems that are substantially the same. One notable exception is the work of Wojtaszczyk, which demonstrates that a slight modification of (2.10) actually obeys a similar result but where the impact of the noise on recovery is actually $\|e\|_2$ rather than the bound ϵ [79]. In other words, (2.10) can be made robust to parameter mismatch in ϵ,

[2] In the event that x is sparse when represented in a basis Ψ, then we must modify (2.10) appropriately and measure the reconstruction error as $\|\widehat{\alpha} - \alpha\|_2$, where $x = \Psi\alpha$. In this case the bound becomes $C_0\epsilon + C_1\sigma_K(\alpha)_1/\sqrt{K}$.

whereas Theorem 3.2 tells us nothing about what happens if $\|e\|_2 > \epsilon$ or if $\|e\|_2 \ll \epsilon$. It remains an open question whether (2.10) itself satisfies this property.

3.3.2 Greedy algorithms

What is perhaps even more surprising is that the RIP is also sufficient for many of the greedy algorithms described in Sections 2.4.2 and 2.4.3 to satisfy results similar to Theorem 3.2. We will have much more to say about the cases of OMP and ROMP in Chapter 6, but as a representative example we provide a modified statement of Theorem 2.2 of [44] on the performance of CoSaMP.

Theorem 3.3 (Needell-Tropp [44])**.** *Suppose that Φ satisfies the RIP of order $4K$ with isometry constant $\delta_{4K} < 0.1$. Given measurements of the form $y = \Phi x + e$, then after $O(K)$ iterations, CoSaMP produces an estimate $\widehat{x}$ satisfying*

$$\|\widehat{x} - x\|_2 \leq C_0 \|e\|_2 + C_1 \frac{\sigma_K(x)_1}{\sqrt{K}}, \tag{3.5}$$

for some constants C_0, C_1 depending only on δ_{4K} and the number of iterations performed.

Note that IHT satisfies a similar result, as well as the SP and DThresh algorithms. Of course, while all of these algorithms have similar performance guarantees, the constants vary widely, and practical performance depends as much on the details of the implementation and usage of the appropriate a "tricks of the trade" as anything else. In general, there is no clear winner among these algorithms at present.

3.4 Measurement Bounds

Finally, we consider how many measurements are *necessary* to achieve the RIP. We first focus on the constant δ.

Theorem 3.4. *Let Φ be an $M \times N$ matrix that satisfies the RIP of order K with constant $\delta_K \in (0,1)$. Then*

$$M \geq \frac{1}{2}\left(\frac{K-1}{\delta_K} - K\right). \tag{3.6}$$

Proof. We begin with the observation that if $K' \geq M+1$, then $\delta_{K'} \geq 1$. This follows from the fact that if $K' \geq M+1$, then any $M \times K'$ submatrix of Φ will have a nontrivial null space, and hence the lower bound of (3.2) must be zero. Thus, from Lemma 3.1, we have that if $K' = c\left\lfloor \frac{K}{2} \right\rfloor \geq M+1$ for some integer c, then $1 < c\,\delta_K$. Hence, for any integer c satisfying

$$c \geq \frac{M+1}{\left\lfloor \frac{K}{2} \right\rfloor}$$

we have that

$$c > \frac{1}{\delta_K}.$$

This implies that

$$\begin{aligned}
\frac{1}{\delta_K} &< \left\lceil \frac{M+1}{\left\lfloor \frac{K}{2} \right\rfloor} \right\rceil \\
&\leq \left\lceil \frac{M+1}{\frac{K}{2} - \frac{1}{2}} \right\rceil = \left\lceil \frac{2M+2}{K-1} \right\rceil \\
&\leq \frac{2M+2}{K-1} + \left(1 - \frac{1}{K-1}\right) = \frac{2M+K}{K-1}.
\end{aligned}$$

This reduces to yield the desired result. □

This bound tells us that if we fix the desired RIP order K, as we decrease the RIP constant δ_K, the required number of measurements increases at a rate that is *at least* proportional to $1/\delta_K$.

If we now ignore the impact of δ and focus only on the dimensions of the problem

(N, M, and K) we can establish another lower bound. We first provide a preliminary lemma that we will need in the proof of the main theorem. The lemma is a direct consequence of Lemma 3.1 of [80], which is a well-known "folk theorem" from coding theory.

Lemma 3.2. *Let K and N satisfying $K < N$ be given. For any $\epsilon \in (0, 1 - K/N)$, there exists a set X consisting of length-N binary vectors, each with exactly K ones, such that for any $x, y \in X$,*

$$\sqrt{2K\epsilon} \leq \|x - y\|_2 \leq \sqrt{2K}, \tag{3.7}$$

and

$$\log |X| > \left(1 - H_{\frac{N}{K}}(\epsilon)\right) K \log\left(\frac{N}{K}\right), \tag{3.8}$$

where H_q is the q-ary entropy function

$$H_q(\epsilon) = -\epsilon \log_q\left(\frac{\epsilon}{q-1}\right) - (1-\epsilon)\log_q(1-\epsilon).$$

Using this lemma, we can establish the following lower bound.

Theorem 3.5. *Let Φ be an $M \times N$ matrix that satisfies the RIP of order $K \leq N/2$ with constant $\delta \in (0, 1)$. Then*

$$M \geq C_\delta K \log\left(\frac{N}{K}\right), \tag{3.9}$$

where $C_\delta < 1$ is a constant depending only on δ.

Proof. Since Φ satisfies (3.2), we also have that

$$\sqrt{1-\delta}\|x - y\|_2 \leq \|\Phi x - \Phi y\|_2 \leq \sqrt{1+\delta}\|x - y\|_2 \tag{3.10}$$

for all $x, y \in \Sigma_K$, since $x - y \in \Sigma_{2K}$. Now consider the set of points X in Lemma 3.2

for some $\epsilon < 1/2$. By construction, we have that

$$\sqrt{2K\epsilon} \leq \|x - y\|_2 \leq \sqrt{2K}, \tag{3.11}$$

and since $X \subset \Sigma_K$, we can combine (3.10) and (3.11) to obtain

$$\sqrt{2K\epsilon(1-\delta)} \leq \|\Phi x - \Phi y\|_2 \leq \sqrt{2K(1+\delta)}$$

for all $x, y \in X$. From the lower bound we can say that for any pair of points $x, y \in X$, if we center balls of radius $\sqrt{2K\epsilon(1-\delta)}/2 = \sqrt{K\epsilon(1-\delta)/2}$ at Φx and Φy, then these balls will be disjoint. In turn, the upper bound tells us that the maximum difference between the centers of any pair of balls is $\sqrt{2K(1+\delta)}$, which means that the entire set of balls is itself contained within a larger ball of radius $\sqrt{2K(1+\delta)} + \sqrt{K\epsilon(1-\delta)/2}$. This implies that

$$\mathrm{Vol}\left(B^M\left(\sqrt{2K(1+\delta)} + \sqrt{\frac{K\epsilon(1-\delta)}{2}}\right)\right) \geq |X| \cdot \mathrm{Vol}\left(B^M\left(\sqrt{\frac{K\epsilon(1-\delta)}{2}}\right)\right),$$

where $\mathrm{Vol}\left(B^M(r)\right)$ denotes the volume of a ball of radius r in $\mathbb{R}^M$. This can be expressed equivalently as

$$\left(\sqrt{2K(1+\delta)} + \sqrt{\frac{K\epsilon(1-\delta)}{2}}\right)^M \geq |X| \cdot \left(\sqrt{\frac{K\epsilon(1-\delta)}{2}}\right)^M,$$

or equivalently,

$$\left(\sqrt{\frac{2(1+\delta)}{\epsilon(1-\delta)}} + 1\right)^M \geq |X|.$$

This reduces to

$$M \geq \frac{\log|X|}{\log\left(2\sqrt{\frac{2(1+\delta)}{\epsilon(1-\delta)}} + 1\right)}.$$

Applying the bound for $|X|$ from (3.2) of Lemma 3.2 and setting $\epsilon = 1/4$, we obtain

$$M \geq \frac{1 - H_{\frac{N}{K}}\left(\frac{1}{4}\right)}{\log\left(\sqrt{\frac{1+\delta}{1-\delta}} + 1\right)} K \log\left(\frac{N}{K}\right).$$

Observing that H_q is monotonically decreasing as a function of q, we can replace $H_{\frac{N}{K}}$ with H_2, which establishes the theorem with

$$C_\delta \approx \frac{0.18}{\log\left(\sqrt{\frac{1+\delta}{1-\delta}} + 1\right)}.$$

□

Note that there is nothing special about the requirement that $K \leq N/2$; this choice was made only to simplify the argument. We have made no effort to optimize the constants, but it is worth noting that they are already quite reasonable. For example, for the case of $\delta = 1/4$ we have $C_\delta \approx 0.5$.

We also note that Theorem 3.5 agrees with the lower bounds that are implied by the work of Garnaev, Gluskin, and Kashin on n-widths [81, 82]. Specifically, in the 1970's they calculated bounds for the various n-widths for certain ℓ_p balls in $\mathbb{R}^N$, and it is possible to relate these widths to the best possible performance of sparse recovery algorithms. Essentially, the fact that the RIP is sufficient for sparse recovery algorithms to achieve a certain level of performance means that if one were able to obtain a matrix satisfying the RIP with fewer measurements than in (3.9), then this would contradict the existing results on n-widths, and thus we cannot do better than (3.9). See [83–85] for further details. In comparison to these previous results, Theorem 3.5 has an appealing simplicity.

Chapter 4

Random Measurements

We now turn to the problem of generating matrices Φ that satisfy the RIP. While our first instinct might be to develop an explicit procedure for designing such a matrix, we will consider a radically different approach — we will instead pick our matrix Φ *at random*. We will construct our random matrices as follows: given M and N, generate random matrices Φ by choosing the entries ϕ_{ij} as independent realizations from a random distribution. In this chapter,[1] we show that under suitable conditions on this distribution, the required number of measurements is within a constant factor of the lower bound established in Theorem 3.5. Perhaps surprisingly, to date there exist no deterministic constructions of Φ which attain this bound. Moreover, we also show that these same techniques can be applied to show that random matrices produce stable embeddings of other signal models beyond sparsity.

4.1 Sub-Gaussian Random Variables

A number of distributions, notably Gaussian and Bernoulli, are known to satisfy certain concentration of measure inequalities. We further analyze this phenomenon in

[1]Section 4.3 consists of work that was done in collaboration with Richard G. Baraniuk, Ronald A. DeVore, and Michael B. Wakin.

Section 4.2 by considering the more general class of sub-Gaussian distributions [86].

Definition 4.1. A random variable X is called *sub-Gaussian* if there exists a constant $c > 0$ such that

$$\mathbb{E}\left(\exp(Xt)\right) \leq \exp(c^2 t^2/2) \tag{4.1}$$

holds for all $t \in \mathbb{R}$. We use the notation $X \sim \mathrm{Sub}(c^2)$ to denote that X satisfies (4.1).

The function $\mathbb{E}\left(\exp(Xt)\right)$ is the *moment-generating function* of X, while the upper bound in (4.1) is the moment-generating function of a Gaussian random variable. Thus, a sub-Gaussian distribution is one whose moment-generating function is bounded by that of a Gaussian. There are a tremendous number of sub-Gaussian distributions, but there are two particularly important examples:

- If $X \sim \mathcal{N}(0, \sigma^2)$, i.e., X is a zero-mean Gaussian random variable with variance σ^2, then $X \sim \mathrm{Sub}(\sigma^2)$. Indeed, as mentioned above, the moment-generating function of a Gaussian is given by $\mathbb{E}\left(\exp(Xt)\right) = \exp(\sigma^2 t^2/2)$, and thus (4.1) is trivially satisfied.

- If X is a zero-mean, bounded random variable, i.e., one for which there exists a constant B such that $|X| \leq B$ with probability 1, then $X \sim \mathrm{Sub}(B^2)$.

A common way to characterize sub-Gaussian random variables is through analyzing their moments. We consider only the mean and variance in the following elementary lemma, proven in [86].

Lemma 4.1 (Buldygin-Kozachenko [86])**.** *If $X \sim \mathrm{Sub}(c^2)$ then,*

$$\mathbb{E}(X) = 0 \tag{4.2}$$

and

$$\mathbb{E}(X^2) \leq c^2. \tag{4.3}$$

Lemma 4.1 shows that if $X \sim \text{Sub}(c^2)$ then $\mathbb{E}(X^2) \leq c^2$. In some cases it will be useful to consider a more restrictive class of random variables for which this inequality becomes an equality.

Definition 4.2. A random variable X is called *strictly sub-Gaussian* if $X \sim \text{Sub}(\sigma^2)$ where $\sigma^2 = \mathbb{E}(X^2)$, i.e., the inequality

$$\mathbb{E}\left(\exp(Xt)\right) \leq \exp(\sigma^2 t^2/2) \tag{4.4}$$

holds for all $t \in \mathbb{R}$. To denote that X is strictly sub-Gaussian with variance σ^2, we will use the notation $X \sim \text{SSub}(\sigma^2)$.

Examples of strictly sub-Gaussian distributions include:

- If $X \sim \mathcal{N}(0, \sigma^2)$, then $X \sim \text{SSub}(\sigma^2)$.
- If $X \sim \text{U}(-1, 1)$, i.e., X is uniformly distributed on the interval $[-1, 1]$, then $X \sim \text{SSub}(1/3)$.
- Now consider the random variable with distribution such that

$$\mathbb{P}(X = 1) = \mathbb{P}(X = -1) = \frac{1 - s}{2}, \quad \mathbb{P}(X = 0) = s, \quad s \in [0, 1).$$

  For any $s \in [0, 2/3]$, $X \sim \text{SSub}(1 - s)$. For $s \in (2/3, 1)$, X is not strictly sub-Gaussian.

We now provide an equivalent characterization for sub-Gaussian and strictly sub-Gaussian random variables, proven in [86], that illustrates their concentration of measure behavior.

Theorem 4.1 (Buldygin-Kozachenko [86])**.** *A random variable $X \sim \mathrm{Sub}(c^2)$ if and only if there exists a $t_0 \geq 0$ and a constant $a \geq 0$ such that*

$$\mathbb{P}(|X| \geq t) \leq 2 \exp\left(-\frac{t^2}{2a^2}\right) \tag{4.5}$$

for all $t \geq t_0$. Moreover, if $X \sim \mathrm{SSub}(\sigma^2)$, then (4.5) holds for all $t > 0$ with $a = \sigma$.

Finally, sub-Gaussian distributions also satisfy one of the fundamental properties of a Gaussian distribution: the sum of two sub-Gaussian random variables is itself a sub-Gaussian random variable. This result is established in more generality in the following lemma.

Lemma 4.2. *Suppose that $X = [X_1, X_2, \ldots, X_N]$, where each X_i is independent and identically distributed (i.i.d.) with $X_i \sim \mathrm{Sub}(c^2)$. Then for any $\alpha \in \mathbb{R}^N$, $\langle X, \alpha \rangle \sim \mathrm{Sub}\left(c^2 \|\alpha\|_2^2\right)$. Similarly, if each $X_i \sim \mathrm{SSub}(\sigma^2)$, then for any $\alpha \in \mathbb{R}^N$, $\langle X, \alpha \rangle \sim \mathrm{SSub}\left(\sigma^2 \|\alpha\|_2^2\right)$.*

Proof. Since the X_i are i.i.d., the joint distribution factors and simplifies as:

$$\begin{aligned}
\mathbb{E}\left(\exp\left(t \sum_{i=1}^{N} \alpha_i X_i\right)\right) &= \mathbb{E}\left(\prod_{i=1}^{N} \exp\left(t \alpha_i X_i\right)\right) \\
&= \prod_{i=1}^{N} \mathbb{E}\left(\exp\left(t \alpha_i X_i\right)\right) \\
&\leq \prod_{i=1}^{N} \exp\left(c^2 (\alpha_i t)^2 / 2\right) \\
&= \exp\left(\left(\sum_{i=1}^{N} \alpha_i^2\right) c^2 t^2 / 2\right).
\end{aligned}$$

In the case where the X_i are strictly sub-Gaussian, the result follows by setting $c^2 = \sigma^2$ and observing that $\mathbb{E}\left(\langle X, \alpha \rangle^2\right) = \sigma^2 \|\alpha\|_2^2$. □

4.2 Sub-Gaussian Matrices and Concentration of Measure

Sub-Gaussian distributions have a close relationship to the concentration of measure phenomenon [87]. To illustrate this, we note that Lemma 4.2 allows us to apply Theorem 4.1 to obtain deviation bounds for weighted sums of sub-Gaussian random variables. For our purposes, however, it will be more interesting to study the *norm* of a vector of sub-Gaussian random variables. In particular, if X is a vector where each X_i is i.i.d. with $X_i \sim \text{Sub}(c)$, we would like to know how $\|X\|_2$ deviates from its mean.

In order to establish the result, we will make use of Markov's inequality for nonnegative random variables.

Lemma 4.3 (Markov)**.** *For any nonnegative random variable X and $t > 0$,*

$$\mathbb{P}(X \geq t) \leq \frac{\mathbb{E}(X)}{t}.$$

Proof. Let $f(x)$ denote the probability density function for X.

$$\mathbb{E}(X) = \int_0^\infty x f(x)\, dx \geq \int_t^\infty x f(x)\, dx \geq \int_t^\infty t f(x)\, dx = t\mathbb{P}(X \geq t).$$

□

In addition, we will require the following bound on the exponential moment of a sub-Gaussian random variable.

Lemma 4.4. *Suppose $X \sim \text{Sub}(c^2)$. Then*

$$\mathbb{E}\left(\exp(\lambda X^2/2c^2)\right) \leq \frac{1}{\sqrt{1-\lambda}}, \tag{4.6}$$

for any $\lambda \in [0, 1)$.

Proof. First, observe that in the case where $\lambda = 0$, the lemma holds trivially. Thus, suppose that $\lambda \in (0, 1)$. Let $f(x)$ denote the probability density function for X. Since X is sub-Gaussian, we have by definition that

$$\int_{-\infty}^{\infty} \exp(tx) f(x)\, dx \leq \exp(c^2 t^2/2)$$

for any $t \in \mathbb{R}$. If we multiply by $\exp(-c^2 t^2/2\lambda)$, we obtain

$$\int_{-\infty}^{\infty} \exp(tx - c^2 t^2/2\lambda) f(x)\, dx \leq \exp(c^2 t^2(\lambda - 1)/2\lambda).$$

Now, integrating both sides with respect to t, we obtain

$$\int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} \exp(tx - c^2 t^2/2\lambda)\, dt \right) f(x)\, dx \leq \int_{-\infty}^{\infty} \exp(c^2 t^2(\lambda - 1)/2\lambda)\, dt,$$

which reduces to

$$\frac{1}{c}\sqrt{2\pi\lambda} \int_{-\infty}^{\infty} \exp(\lambda x^2/2c^2) f(x)\, dx \leq \frac{1}{c}\sqrt{\frac{2\pi\lambda}{1-\lambda}}.$$

This simplifies to prove the lemma. □

We now state our main theorem, which generalizes the results of [88] and uses substantially the same proof technique.

Theorem 4.2. *Suppose that $X = [X_1, X_2, \ldots, X_M]$, where each X_i is i.i.d. with $X_i \sim \mathrm{Sub}(c^2)$ and $\mathbb{E}(X_i^2) = \sigma^2$. Then*

$$\mathbb{E}\left(\|X\|_2^2\right) = M\sigma^2. \tag{4.7}$$

Moreover, for any $\alpha \in (0, 1)$ and for any $\beta \in [c^2/\sigma^2, \beta_{\max}]$, there exists a constant

$C^* \geq 4$ depending only on $\beta_{\max}$ and the ratio σ^2/c^2 such that

$$\mathbb{P}\left(\|X\|_2^2 \leq \alpha M\sigma^2\right) \leq \exp\left(-M(1-\alpha)^2/C^*\right) \tag{4.8}$$

and

$$\mathbb{P}\left(\|X\|_2^2 \geq \beta M\sigma^2\right) \leq \exp\left(-M(\beta-1)^2/C^*\right). \tag{4.9}$$

Proof. Since the X_i are independent, we obtain

$$\mathbb{E}\left(\|X\|_2^2\right) = \sum_{i=1}^{M} \mathbb{E}\left(X_i^2\right) = \sum_{i=1}^{M} \sigma^2 = M\sigma^2$$

and hence (4.7) holds. We now turn to (4.8) and (4.9). Let us first consider (4.9). We begin by applying Markov's inequality:

$$\begin{aligned}
\mathbb{P}\left(\|X\|_2^2 \geq \beta M\sigma^2\right) &= \mathbb{P}\left(\exp(\lambda\|X\|_2^2) \geq \exp\left(\lambda\beta M\sigma^2\right)\right) \\
&\leq \frac{\mathbb{E}\left(\exp(\lambda\|X\|_2^2)\right)}{\exp\left(\lambda\beta M\sigma^2\right)} \\
&= \frac{\prod_{i=1}^{M} \mathbb{E}\left(\exp(\lambda X_i^2)\right)}{\exp\left(\lambda\beta M\sigma^2\right)}.
\end{aligned}$$

Since $X_i \sim \mathrm{Sub}(c^2)$, we have from Lemma 4.4 that

$$\mathbb{E}\left(\exp(\lambda X_i^2)\right) = \mathbb{E}\left(\exp(2c^2\lambda X_i^2/2c^2)\right) \leq \frac{1}{\sqrt{1-2c^2\lambda}}.$$

Thus,

$$\prod_{i=1}^{M} \mathbb{E}\left(\exp\left(\lambda X_i^2\right)\right) \leq \left(\frac{1}{1-2c^2\lambda}\right)^{M/2}$$

and hence

$$\mathbb{P}\left(\|X\|_2^2 \geq \beta M\sigma^2\right) \leq \left(\frac{\exp\left(-2\lambda\beta\sigma^2\right)}{1-2c^2\lambda}\right)^{M/2}.$$

By setting the derivative to zero and solving for λ, one can show that the optimal λ

is

$$\lambda = \frac{\beta\sigma^2 - c^2}{2c^2\sigma^2(1+\beta)}.$$

Plugging this in we obtain

$$\mathbb{P}\left(\|X\|_2^2 \geq \beta M\sigma^2\right) \leq \left(\beta\frac{\sigma^2}{c^2}\exp\left(1-\beta\frac{\sigma^2}{c^2}\right)\right)^{M/2}. \tag{4.10}$$

Similarly,

$$\mathbb{P}\left(\|X\|_2^2 \leq \alpha M\sigma^2\right) \leq \left(\alpha\frac{\sigma^2}{c^2}\exp\left(1-\alpha\frac{\sigma^2}{c^2}\right)\right)^{M/2}. \tag{4.11}$$

In order to combine and simplify these inequalities, we will make use of the fact that if we define

$$C^* = \max\left(4, 2\frac{(\beta_{\max}\sigma^2/c-1)^2}{(\beta_{\max}\sigma^2/c-1)-\log(\beta_{\max}\sigma^2/c)}\right)$$

then we have that for any $\gamma \in [0, \beta_{\max}\sigma^2/c]$ we have the bound

$$\log(\gamma) \leq (\gamma-1) - \frac{2(\gamma-1)^2}{C^*}, \tag{4.12}$$

and hence

$$\gamma \leq \exp\left((\gamma-1) - \frac{2(\gamma-1)^2}{C^*}\right).$$

By setting $\gamma = \alpha\sigma^2/c^2$, (4.11) reduces to yield (4.8). Similarly, setting $\gamma = \beta\sigma^2/c^2$ establishes (4.9). □

This result tells us that given a vector with entries drawn according to a sub-Gaussian distribution, we can expect the norm of the vector to concentrate around its expected value of $M\sigma^2$ with exponentially high probability as M grows. Note, however, that the range of allowable choices for β in (4.9) is limited to $\beta \geq c^2/\sigma^2 \geq 1$. Thus, for a general sub-Gaussian distribution, we may be unable to achieve an arbitrarily tight concentration. However, recall that for strictly sub-Gaussian distributions we have that $c^2 = \sigma^2$, in which there is no such restriction. Moreover, for

strictly sub-Gaussian distributions we also have the following useful corollary.[2]

Corollary 4.1. *Suppose that* $X = [X_1, X_2, \ldots, X_M]$*, where each* X_i *is i.i.d. with* $X_i \sim \mathrm{SSub}(\sigma^2)$*. Then*

$$\mathbb{E}\left(\|X\|_2^2\right) = M\sigma^2 \tag{4.13}$$

and for any $\epsilon > 0$*,*

$$\mathbb{P}\left(\left|\|X\|_2^2 - M\sigma^2\right| \geq \epsilon M\sigma^2\right) \leq 2\exp\left(-\frac{M\epsilon^2}{C^*}\right) \tag{4.14}$$

with $C^* = 2/(1 - \log(2)) \approx 6.52$*.*

Proof. Since each $X_i \sim \mathrm{SSub}(\sigma^2)$, we have that $X_i \sim \mathrm{Sub}(\sigma^2)$ and $\mathbb{E}(X_i^2) = \sigma^2$, in which case we may apply Theorem 4.2 with $\alpha = 1 - \epsilon$ and $\beta = 1 + \epsilon$. This allows us to simplify and combine the bounds in (4.8) and (4.9) to obtain (4.14). The value of C^* follows from the fact that in this case we have that $1 + \epsilon \leq 2$ so that we can set $\beta_{\max} = 2$. □

Finally, from Corollary 4.1 we also have the following additional corollary that we will use in Section 4.3. This result generalizes the main results of [89] to the broader family of general strictly sub-Gaussian distributions via a much simpler proof. Note that the conclusion of this corollary is also essentially the same as Lemma 6.1 of [90].

Corollary 4.2. *Suppose that* Φ *is an* $M \times N$ *matrix whose entries* ϕ_{ij} *are i.i.d. with* $\phi_{ij} \sim \mathrm{SSub}(1/M)$*. Let* $Y = \Phi x$ *for* $x \in \mathbb{R}^N$*. Then for any* $\epsilon > 0$*, and any* $x \in \mathbb{R}^N$*,*

$$\mathbb{E}\left(\|Y\|_2^2\right) = \|x\|_2^2 \tag{4.15}$$

[2]Corollary 4.1 exploits the strictness in the strictly sub-Gaussian distribution twice — first to ensure that $\beta \in (1, 2]$ is an admissible range for β and then to simplify the computation of C^*. One could easily establish a different version of this corollary for non-strictly sub-Gaussian vectors but for which we consider a more restricted range of ϵ provided that $c^2/\sigma^2 < 2$. However, since most of the distributions of interest in this thesis are indeed strictly sub-Gaussian, we do not pursue this route. Note also that if one is interested only in the case where ϵ is very small, there is considerable room for improvement in the constant C^*.

and

$$\mathbb{P}\left(\left|\|Y\|_2^2 - \|x\|_2^2\right| \geq \epsilon\|x\|_2^2\right) \leq 2\exp\left(-\frac{M\epsilon^2}{C^*}\right) \tag{4.16}$$

with $C^* = 2/(1-\log(2)) \approx 6.52$.

Proof. Let ϕ_i denote the i^{th} row of Φ. Observe that if Y_i denotes the first element of Y, then $Y_i = \langle \phi_i, x \rangle$, and thus by Lemma 4.2, $Y_i \sim \text{SSub}\left(\|x\|_2^2/M\right)$. Applying Corollary 4.1 to the M-dimensional random vector Y, we obtain (4.16). □

4.3 Sub-Gaussian Matrices and the RIP

We now show how to exploit the concentration of measure properties of sub-Gaussian distributions to provide a simple proof that sub-Gaussian matrices satisfy the RIP. We begin by observing that if all we require is that $\delta_{2K} > 0$, then we may set $M = 2K$ and draw a Φ according to a Gaussian distribution, or indeed any continuous univariate distribution. In this case, with probability 1, any subset of $2K$ columns will be linearly independent, and hence all subsets of $2K$ columns will be bounded below by $1 - \delta_{2K}$ where $\delta_{2K} > 0$. However, suppose we wish to know the constant δ_{2K}. In order to find the value of the constant we must consider all possible $\binom{N}{K}$ K-dimensional subspaces of $\mathbb{R}^N$. From a computational perspective, this is impossible for any realistic values of N and K. Moreover, in light of Theorems 3.4 and 3.5, the actual value of δ_{2K} in this case is likely to be very close to 1. Thus, we focus instead on the problem of achieving the RIP of order $2K$ for a specified constant δ_{2K}.

To ensure that the matrix will satisfy the RIP, we will impose two conditions on the random distribution. First, we require that the distribution is sub-Gaussian. In order to simplify our argument, we will use the simpler results stated in Corollaries 4.1 and 4.2, so our theorem will actually assume that the distribution is strictly sub-Gaussian, although the argument could also be modified to establish a similar result for general sub-Gaussian distributions using Theorem 4.2.

Our second condition is that we require that the distribution yield a matrix that is approximately norm-preserving, which will require that

$$\mathbb{E}(\phi_{ij}^2) = \frac{1}{M}, \tag{4.17}$$

and hence the variance is $1/M$.

We shall now show how the concentration of measure inequality in Corollary 4.2 can be used together with covering arguments to prove the RIP for sub-Gaussian random matrices. Our general approach will be to construct nets of points in each K-dimensional subspace, apply (4.16) to all of these points through a union bound, and then extend the result from our finite set of points to all possible K-dimensional signals. Thus, in order to prove the result, we will require the following upper bound on the number of points required to construct the nets of points. (For an overview of results similar to Lemma 4.5 and of various related concentration of measure results, we refer the reader to the excellent introduction of [91].)

Lemma 4.5. *Let $\epsilon \in (0,1)$ be given. There exists a set of points Q such that $|Q| \leq (3/\epsilon)^K$ and for any $x \in \mathbb{R}^K$ with $\|x\|_2 \leq 1$ there is a point $q \in Q$ satisfying $\|x-q\|_2 \leq \epsilon$.*

Proof. We construct Q in a greedy fashion. We first select an arbitrary point $q_1 \in \mathbb{R}^K$ with $\|q_1\|_2 \leq 1$. We then continue adding points to Q so that at step i we add a point $q_i \in \mathbb{R}^K$ with $\|q_i\|_2 \leq 1$ which satisfies $\|q_i - q_j\|_2 > \epsilon$ for all $j < i$. This continues until we can add no more points (and hence for any $x \in \mathbb{R}^K$ with $\|x\|_2 \leq 1$ there is a point $q \in Q$ satisfying $\|x - q\|_2 \leq \epsilon$.) Now we wish to bound $|Q|$. Observe that if we center balls of radius $\epsilon/2$ at each point in Q, then these balls are disjoint and lie within a ball of radius $1 + \epsilon/2$. Thus, if $B^K(r)$ denotes a ball of radius r in $\mathbb{R}^K$, then

$$|Q| \cdot \text{Vol}\left(B^K(\epsilon/2)\right) \leq \text{Vol}\left(B^K(1+\epsilon/2)\right)$$

and hence

$$\begin{aligned} |Q| &\leq \frac{\mathrm{Vol}\left(B^K(1+\epsilon/2)\right)}{\mathrm{Vol}\left(B^K(\epsilon/2)\right)} \\ &= \frac{(1+\epsilon/2)^K}{(\epsilon/2)^K} \\ &\leq (3/\epsilon)^K . \end{aligned}$$

□

We now turn to our main theorem.

Theorem 4.3. *Fix $\delta \in (0,1)$. Let Φ be an $M \times N$ random matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$. If*

$$M \geq C_1 K \log\left(\frac{N}{K}\right), \tag{4.18}$$

then Φ satisfies the RIP of order K with the prescribed δ with probability exceeding $1 - 2e^{-C_2 M}$, where C_1 is arbitrary and $C_2 = \delta^2/2C^* - \log(42e/\delta)/C_1$.

Proof. First note that it is enough to prove (3.2) in the case $\|x\|_2 = 1$, since Φ is linear. Next, fix an index set $T \subset \{1, 2, \ldots, N\}$ with $|T| = K$, and let X_T denote the K-dimensional subspace spanned by the columns of Φ_T. We choose a finite set of points Q_T such that $Q_T \subseteq X_T$, $\|q\|_2 \leq 1$ for all $q \in Q_T$, and for all $x \in X_T$ with $\|x\|_2 \leq 1$ we have

$$\min_{q \in Q_T} \|x - q\|_2 \leq \delta/14. \tag{4.19}$$

From Lemma 4.5, we can choose such a set Q_T with $|Q_T| \leq (42/\delta)^K$. We then repeat this process for each possible index set T, and collect all the sets Q_T together:

$$Q = \bigcup_{T:|T|=K} Q_T. \tag{4.20}$$

There are $\binom{N}{K}$ possible index sets T. We can bound this number by

$$\binom{N}{K} = \frac{N(N-1)(N-2)\cdots(N-K+1)}{K!} \leq \frac{N^K}{K!} \leq \left(\frac{eN}{K}\right)^K,$$

where the last inequality follows since from Sterling's approximation we have $K! \geq (K/e)^K$. Hence $|Q| \leq (42eN/\delta K)^K$. Since the entries of Φ are drawn according to a strictly sub-Gaussian distribution, from Corollary 4.2 we have (4.16). We next use the union bound to apply (4.16) to this set of points with $\epsilon = \delta/\sqrt{2}$, with the result that, with probability exceeding

$$1 - 2(42eN/\delta K)^K \, e^{-M\delta^2/2C^*}, \tag{4.21}$$

we have

$$(1-\delta/\sqrt{2})\|q\|_2^2 \leq \|\Phi q\|_2^2 \leq (1+\delta/\sqrt{2})\|q\|_2^2, \quad \text{for all } q \in Q. \tag{4.22}$$

We observe that if M satisfies (4.18) then

$$\log\left(\frac{42eN}{\delta K}\right)^K \leq K \log\left(\frac{N}{K}\right) \log\left(\frac{42e}{\delta}\right) \leq \frac{M \log(42e/\delta)}{C_1}$$

and thus (4.21) exceeds $1 - 2e^{-C_2 M}$ as desired.

We now define A as the smallest number such that

$$\|\Phi x\|_2 \leq \sqrt{1+A}\|x\|_2, \quad \text{for all } x \in \Sigma_K, \ \|x\|_2 \leq 1. \tag{4.23}$$

Our goal is to show that $A \leq \delta$. For this, we recall that for any $x \in \Sigma_K$ with $\|x\|_2 \leq 1$, we can pick a $q \in Q$ such that $\|x - q\|_2 \leq \delta/14$ and such that $x - q \in \Sigma_K$ (since if

$x \in X_T$, we can pick $q \in Q_T \subset X_T$ satisfying $\|x - q\|_2 \leq \delta/14$). In this case we have

$$\|\Phi x\|_2 \leq \|\Phi q\|_2 + \|\Phi(x-q)\|_2 \leq \sqrt{1+\delta/\sqrt{2}} + \sqrt{1+A} \cdot \delta/14. \tag{4.24}$$

Since by definition A is the smallest number for which (4.23) holds, we obtain $\sqrt{1+A} \leq \sqrt{1+\delta/\sqrt{2}} + \sqrt{1+A} \cdot \delta/14$. Therefore

$$\sqrt{1+A} \leq \frac{\sqrt{1+\delta/\sqrt{2}}}{1-\delta/14} \leq \sqrt{1+\delta},$$

as desired. We have proved the upper inequality in (3.2). The lower inequality follows from this since

$$\|\Phi x\|_2 \geq \|\Phi q\|_2 - \|\Phi(x-q)\|_2 \geq \sqrt{1-\delta/\sqrt{2}} - \sqrt{1+\delta} \cdot \delta/14 \geq \sqrt{1-\delta}, \tag{4.25}$$

which completes the proof. □

Above we prove above that the RIP holds with high probability when the matrix Φ is drawn according to a strictly sub-Gaussian distribution. However, we are often interested in signals that are sparse or compressible in some orthonormal basis $\Psi \neq I$, in which case we would like the matrix $\Phi\Psi$ to satisfy the RIP. In this setting it is easy to see that by choosing our net of points in the K-dimensional subspaces spanned by sets of K columns of Ψ, Theorem 4.3 will establish the RIP for $\Phi\Psi$ for Φ again drawn from a sub-Gaussian distribution. This *universality* of Φ with respect to the sparsity-inducing basis is an attractive property that was initially observed for the Gaussian distribution (based on symmetry arguments), but we can now see is a property of more general sub-Gaussian distributions. Indeed, it follows that with high probability such a Φ will simultaneously satisfy the RIP with respect to an exponential number of fixed bases.

4.4 Beyond Sparsity

We now briefly discuss the more general question of how to construct linear mappings Φ that provide stable embeddings for a few important sets besides Σ_K. Specifically, we will examine when we can find Φ satisfying

$$(1-\delta)\|x-y\|_2^2 \leq \|\Phi x - \Phi y\|_2^2 \leq (1+\delta)\|x-y\|_2^2 \tag{4.26}$$

for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ for sets $\mathcal{X}, \mathcal{Y}$ other than $\mathcal{X} = \mathcal{Y} = \Sigma_K$.[3] To make the dependence of (4.26) on δ clear, we will also sometimes use the term δ-stable embedding.

We start with the simple case where we desire a stable embedding of $\mathcal{X} = \mathcal{U}$ and $\mathcal{Y} = \mathcal{V}$ where $\mathcal{U} = \{u_i\}_{i=1}^{|\mathcal{U}|}$ and $\mathcal{V} = \{v_i\}_{i=1}^{|\mathcal{V}|}$ are finite sets of points in $\mathbb{R}^N$. Note that in the case where $\mathcal{U} = \mathcal{V}$, this result is equivalent to the well-known Johnson-Lindenstrauss (JL) lemma [88, 89, 92].

Lemma 4.6. *Let $\mathcal{U}$ be a set of points in $\mathbb{R}^N$. Fix $\delta \in (0,1)$. Let Φ be an $M \times N$ random matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$. If*

$$M \geq C_1 \log(|\mathcal{U}||\mathcal{V}|), \tag{4.27}$$

then Φ is a δ-stable embedding of $(\mathcal{U}, \mathcal{V})$ for the prescribed δ with probability exceeding $1 - 2e^{-C_2 M}$, where C_1 is arbitrary and $C_2 = \delta^2/C^* - 1/C_1$.

Proof. To prove the result we apply (4.16) from Corollary 4.2 to the $|\mathcal{U}| \cdot |\mathcal{V}|$ vectors corresponding to all possible $u_i - v_j$. By applying the union bound, we obtain that the

[3] Note that we are slightly refining our notion of a stable embedding compared to (3.1) by considering the case where x and y come from different sets. This formulation has some advantages that will become apparent in later chapters. For now, it simplifies some of our discussion in this section, since in some cases we will be interested in simply preserving the norms of individual elements in our set rather than distances between pairs of elements in our set. Specifically, we will consider stable embeddings of general sets $\mathcal{X}$ with $\mathcal{Y} = \{0\}$. Note that in general, if Φ is a stable embedding of $(\mathcal{X}, \mathcal{Y})$, this is equivalent to it being a stable embedding of $(\widetilde{\mathcal{X}}, \{0\})$ where $\widetilde{\mathcal{X}} = \{x - y : x \in \mathcal{X}, y \in \mathcal{Y}\}$. This formulation can sometimes be more convenient.

probability of (4.26) not holding is bounded above by $2|\mathcal{U}| \cdot |\mathcal{V}| e^{-M\delta^2/C^*}$. By bounding

$$|\mathcal{U}| \cdot |\mathcal{V}| = e^{\log(|\mathcal{U}| \cdot |\mathcal{V}|)} \leq e^{M/C_1}$$

we obtain the desired result. □

We next note that the technique we used in the proof of Theorem 4.3 provides us with two immediate corollaries. Briefly, recall that the proof of Theorem 4.3 essentially consisted of constructing an ϵ-net of points with ϵ sufficiently small, applying (4.16) to argue that the structure of the ϵ-net was preserved by Φ, and then extending this result to all of Σ_K. We constructed the ϵ-net by picking $(3/\epsilon)^K$ points from each subspace, and then repeating this for all $\binom{N}{K}$ subspaces. If instead we consider only a single subspace, or some subset of subspaces, the same technique yields the following corollaries, which we state without proof.

Corollary 4.3. *Suppose that $\mathcal{X}$ is a K-dimensional subspace of $\mathbb{R}^N$. Fix $\delta \in (0,1)$. Let Φ be an $M \times N$ random matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$. If*

$$M \geq C_1 K, \tag{4.28}$$

then Φ is a δ-stable embedding of $(\mathcal{X}, \{0\})$ (or equivalently, $(\mathcal{X}, \mathcal{X})$) for the prescribed δ with probability exceeding $1 - 2e^{-C_2 M}$, where C_1 is arbitrary and $C_2 = \delta^2/2C^* - \log(42/\delta)/C_1$.

Corollary 4.4. *Suppose that $\mathcal{X} = \cup_{i=1}^{L} \mathcal{X}_i$ is a union of L different K-dimensional subspaces of $\mathbb{R}^N$. Fix $\delta \in (0,1)$. Let Φ be an $M \times N$ random matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$. If*

$$M \geq C_1 K \log(L), \tag{4.29}$$

then Φ is a δ-stable embedding of $(\mathcal{X}, \{0\})$ for the prescribed δ with probability exceed-

ing $1 - 2e^{-C_2 M}$*, where* C_1 *is arbitrary and* $C_2 = \delta^2/2C^* - \log(42/\delta)/C_1$.

Note that if $\mathcal{X}$ is a union of L subspaces of dimension K, then the set of all difference vectors between pairs of points in $\mathcal{X}$ is itself a union of $\binom{L}{2} < L^2$ subspaces of dimension $2K$, so that Corollary 4.4 also immediately provides an argument for a stable embedding of $(\mathcal{X}, \mathcal{X})$ with only a slight change in C_2. See the work of [93, 94] for examples of signal models that consist of a union of subspaces where $L \ll \binom{N}{K}$, in which case Corollary 4.4 can offer significant improvement in terms of the required number of measurements compared to standard RIP-based analysis.

Finally, we note that a similar technique has recently been used to demonstrate that random projections also provide a stable embedding of nonlinear manifolds [95]: under certain assumptions on the curvature and volume of a K-dimensional manifold $\mathcal{M} \subset \mathbb{R}^N$, a random sensing matrix with $M = O\left(K \log(N)\right)$ will with high probability provide a stable embedding of $(\mathcal{M}, \mathcal{M})$. Under slightly different assumptions on $\mathcal{M}$, a number of similar embedding results involving random projections have been established [96–98].

4.5 Random Projections

We will see later, especially in Part IV, that it is often useful to consider *random projection* matrices rather than Φ with i.i.d. sub-Gaussian entries. A random projection is simply an orthogonal projection onto a randomly selected subspace. The two constructions are closely related — the main difference being that the matrices described above are not orthogonal projection matrices, which would require that they have orthonormal rows.

However, random projections share the same key properties as the constructions described above. First, we note that if Φ has orthonormal rows spanning a random subspace, then $\Phi\Phi^T = I$, and so $P_{\Phi^T} = \Phi^T\Phi$. It follows that $\|P_{\Phi^T} s\| = \|\Phi^T\Phi s\| =$

$\|\Phi s\|$, and for random orthogonal projections, it is known [88] that $\|P_{\Phi^T} s\| = \|\Phi s\|$ satisfies

$$(1-\delta)\frac{M}{N}\|s\|^2 \leq \|P_{\Phi^T} s\|^2 \leq (1+\delta)\frac{M}{N}\|s\|^2 \tag{4.30}$$

with probability at least $1-2e^{-CM\delta^2}$ for some constant C. This statement is analogous to (4.16) but rescaled to account for the unit-norm rows of Φ. Note also that if Φ is populated with i.i.d. zero-mean Gaussian entries (of any fixed variance), then the orientation of the row space of Φ has random uniform distribution. Thus, $\|P_{\Phi^T} s\|$ for a Gaussian Φ has the same distribution as $\|P_{\Phi^T} s\|$ for a random orthogonal projection. It follows that Gaussian Φ also satisfy (4.30) with probability at least $1 - 2e^{-CM\delta^2}$.

The similarity between (4.30) and (4.16) immediately implies that we can generalize Theorem 4.3, Lemma 4.6, and Corollaries 4.3 and 4.4 to establish δ-stable embedding results for orthogonal projection matrices P_{Φ^T}. It follows that, when Φ is a Gaussian matrix (with entries satisfying (4.17)) or a random orthogonal projection (multiplied by $\sqrt{N/M}$), the number of measurements required to establish a δ-stable embedding for $\sqrt{N/M}P_{\Phi^T}$ on a particular signal family $\mathcal{S}$ is equivalent to the number of measurements required to establish a δ-stable embedding for Φ on $\mathcal{S}$.

4.6 Deterministic Guarantees and Random Matrix Constructions

Throughout this thesis, we state a variety of theorems that begin with the assumption that Φ is a stable embedding of a set or satisfies the RIP and then use this assumption to establish performance guarantees for a particular algorithm. These guarantees are typically completely deterministic and hold for any Φ that is a stable embedding. However, we use random constructions as our main tool for obtaining stable embeddings. Thus, all of our results could be modified to be probabilistic

statements in which we fix M and then argue that with high probability, a random Φ is a stable embedding. Of course, the concept of "high probability" is somewhat arbitrary. However, if we fix this probability of error to be an acceptable constant ρ, then as we increase M, we are able to reduce δ to be arbitrarily close to 0. This will typically improve the accuracy of the guarantees.

As a side comment, it is important to note that in the case where one is able to generate a new Φ before acquiring each new signal x, then it is sometimes possible to drastically reduce the required M. This is because one may be able to potentially eliminate the requirement that Φ is a stable embedding for an entire class of candidate signals x, and instead simply argue that for each x, a new random matrix Φ with M very small preserves the norm of x, which is sufficient in some settings. Thus, if such a probabilistic "for each" guarantee is acceptable, then it is sometimes possible to place *no assumptions* on the signals being sparse, or indeed having any structure at all. This is particularly true for the results discussed in Part IV. However, in the remainder of this thesis we will restrict ourselves to the sort of deterministic guarantees that hold for a class of signals when Φ provides a stable embedding of that class.

Chapter 5

Compressive Measurements in Practice

In this chapter[1] we discuss various strategies for designing systems for acquiring random, compressive measurements of real-world signals. In order to accomplish this, we must address two main challenges. First, note that the suggested design procedure from Chapter 4 is essentially to pick the entries of Φ at random, in a completely unstructured manner. This approach can be potentially problematic in the case where N is very large, as is typically the case in our applications of interest. This is because the matrix Φ must be stored/transmitted along with the measurements in order to be able to recover the original signal. When N is large, the size of the matrix MN can be can become impractically large. Moreover, the recovery algorithms described in Section 2.4 typically must repeatedly apply the matrix Φ, which in the unstructured case will require $O(MN)$ computations. For large N and M, this cost can easily become prohibitively large. To address these challenges, we will draw on the same techniques often used in the data streaming literature [62] and

[1]This chapter provides an overview of collaborations with and independent work by: Richard G. Baraniuk, Dror Baron, Marco F. Duarte, Kevin F. Kelly, Sami Kirolos, Jason N. Laska, Yehia Massoud, Tamer S. Ragheb, Justin K. Romberg, Shriram Sarvotham, Ting Sun, Dharmpal Takhar, John Treichler, Joel A. Tropp, and Michael B. Wakin.

consider pseudorandom matrices (in which case we only need to store/transmit the random seed used to generate Φ rather than the matrix itself) as well as matrices that have considerable amount of structure, admitting efficient or transform-based implementations.

However, there is a second, potentially even more difficult challenge. The acquisition framework developed in Chapters 3 and 4 assumes that the signal to be acquired is a vector in $\mathbb{R}^N$, while in practice we will be interested in designing systems for continuous-time, analog signals or images. In this chapter we will show that in many cases, it is possible to design a system that directly operates on a continuous-time signal to acquire compressive measurements *without* first sampling the signal. We then show that these measurements can be related to an equivalent system that first samples the signal at its Nyquist-rate, and then applies a matrix Φ to the sampled data.

We primarily focus on two signal acquisition architectures: the single-pixel camera and the random demodulator. The single-pixel camera uses a Texas Instruments DMD array and a single light sensor to optically compute inner products between an image and random patterns. By changing these patterns over time, we can build up a collection of random measurements of an image. The random demodulator provides a CS-inspired hardware architecture for acquiring wideband analog signals.

5.1 The Single-Pixel Camera

5.1.1 Architecture

Several hardware architectures have been proposed that enable the acquisition of compressive measurements in an imaging setting [56–58]. We will focus on the so-called *single-pixel camera* [56, 99–102]. The single-pixel camera is an optical computer that sequentially measures the inner products $y[j] = \langle x, \phi_j \rangle$ between an N-pixel

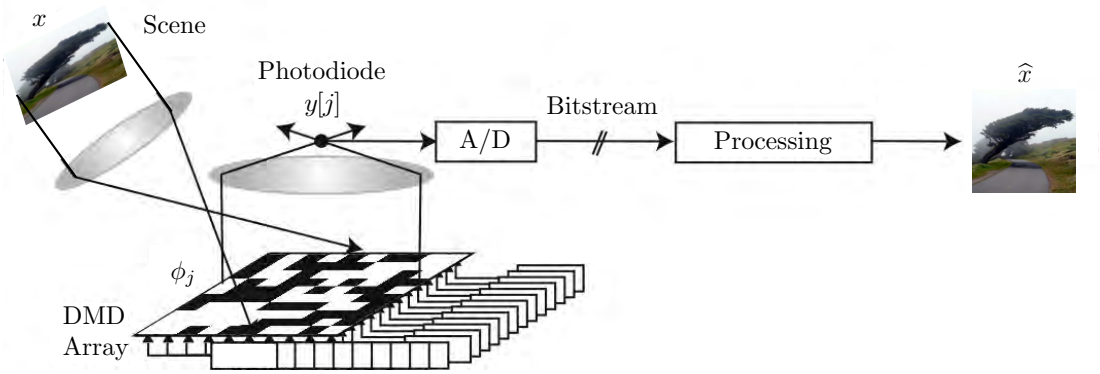

Figure 5.1: Single-pixel camera block diagram. Incident light-field (corresponding to the desired image x) is reflected off a digital micromirror device (DMD) array whose mirror orientations are modulated according to the pseudorandom pattern ϕ_j supplied by a random number generator. Each different mirror pattern produces a voltage at the single photodiode that corresponds to one measurement $y[j]$.

sampled version of the incident light-field from the scene under view (denoted by x) and a set of N-pixel test functions $\{\phi_j\}_{j=1}^{M}$. The architecture is illustrated in Figure 5.1, and an aerial view of the camera in the lab is shown in Figure 5.2. As shown in these figures, the light-field is focused by a lens (Lens 1 in Figure 5.2) not onto a CCD or CMOS sampling array but rather onto a spatial light modulator (SLM). An SLM modulates the intensity of a light beam according to a control signal. A simple example of a transmissive SLM that either passes or blocks parts of the beam is an overhead transparency. Another example is a liquid crystal display (LCD) projector.

The Texas Instruments (TI) digital micromirror device (DMD) is a reflective SLM that selectively redirects parts of the light beam. The DMD consists of an array of bacterium-sized, electrostatically actuated micro-mirrors, where each mirror in the array is suspended above an individual static random access memory (SRAM) cell. Each mirror rotates about a hinge and can be positioned in one of two states (± 10 degrees from horizontal) according to which bit is loaded into the SRAM cell; thus light falling on the DMD can be reflected in two directions depending on the orientation of the mirrors.

Each element of the SLM corresponds to a particular element of ϕ_j (and its cor-

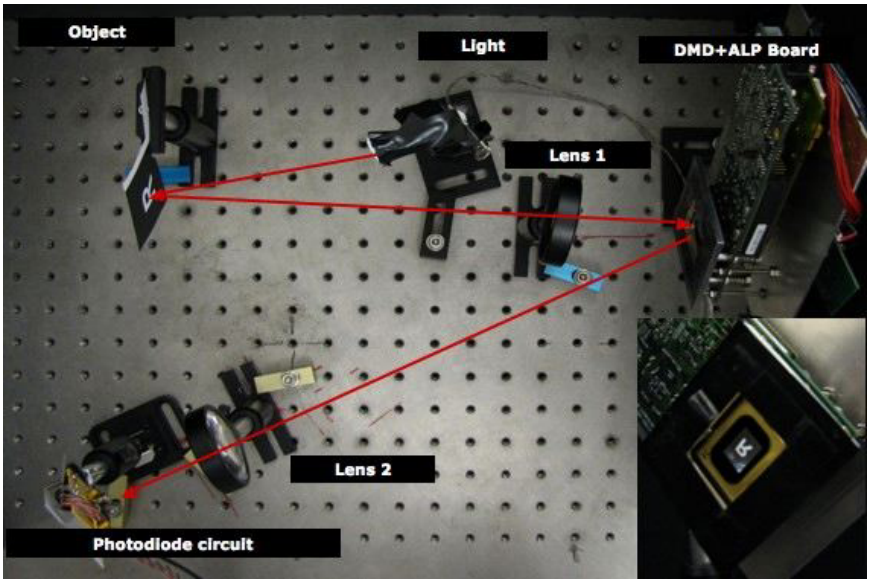

Figure 5.2: Aerial view of the single-pixel camera in the lab.

responding pixel in x). For a given ϕ_j, we can orient the corresponding element of the SLM either towards (corresponding to a 1 at that element of ϕ_j) or away from (corresponding to a 0 at that element of ϕ_j) a second lens (Lens 2 in Figure 5.2). This second lens collects the reflected light and focuses it onto a single photon detector (the single pixel) that integrates the product of x and ϕ_j to compute the measurement $y[j] = \langle x, \phi_j \rangle$ as its output voltage. This voltage is then digitized by an A/D converter. Values of ϕ_j between 0 and 1 can be obtained by dithering the mirrors back and forth during the photodiode integration time. By reshaping x into a column vector and the ϕ_j into row vectors, we can thus model this system as computing the product $y = \Phi x$, where each row of Φ corresponds to a ϕ_j. To compute randomized measurements, we set the mirror orientations ϕ_j randomly using a pseudorandom number generator, measure $y[j]$, and then repeat the process M times to obtain the measurement vector y.

The single-pixel design reduces the required size, complexity, and cost of the photon detector array down to a single unit, which enables the use of exotic detectors that would be impossible in a conventional digital camera. Example detectors include

a photomultiplier tube or an avalanche photodiode for low-light (photon-limited) imaging, a sandwich of several photodiodes sensitive to different light wavelengths for multimodal sensing, a spectrometer for hyperspectral imaging, and so on.

In addition to sensing flexibility, the practical advantages of the single-pixel design include the facts that the quantum efficiency of a photodiode is higher than that of the pixel sensors in a typical CCD or CMOS array and that the fill factor of a DMD can reach 90% whereas that of a CCD/CMOS array is only about 50%. An important advantage to highlight is the fact that each CS measurement receives about $N/2$ times more photons than an average pixel sensor, which significantly reduces image distortion from dark noise and read-out noise.

The single-pixel design falls into the class of multiplex cameras. The baseline standard for multiplexing is classical raster scanning, where the test functions $\{\phi_j\}$ are a sequence of delta functions $\delta[n-j]$ that turn on each mirror in turn. There are substantial advantages to operating in a CS rather than raster scan mode, including fewer total measurements (M for CS rather than N for raster scan) and significantly reduced dark noise. See [56] for a more detailed discussion of these issues.

Figure 5.3 (a) and (b) illustrates a target object (a black-and-white printout of an "R") x and reconstructed image $\widehat{x}$ taken by the single-pixel camera prototype in Figure 5.2 using $N = 256 \times 256$ and $M = N/50$ [56]. Figure 5.3(c) illustrates an $N = 256 \times 256$ color single-pixel photograph of a printout of the Mandrill test image taken under low-light conditions using RGB color filters and a photomultiplier tube with $M = N/10$. In both cases, the images were reconstructed using Total Variation minimization, which is closely related to wavelet coefficient ℓ_1 minimization [15].

5.1.2 Discrete formulation

Since the DMD array is programmable, we can employ arbitrary test functions ϕ_j. However, even when we restrict the ϕ_j to be $\{0, 1\}$-valued, storing these patterns for

(a) (b) (c)

Figure 5.3: Sample image reconstructions from single-pixel camera. (a) 256×256 conventional image of a black-and-white "R". (b) Image reconstructed from $M = 1300$ single-pixel camera measurements ($50\times$ sub-Nyquist). (c) 256×256 pixel color reconstruction of a printout of the Mandrill test image imaged in a low-light setting using a single photomultiplier tube sensor, RGB color filters, and $M = 6500$ random measurements.

large values of N is impractical. Furthermore, as noted above, even pseudorandom Φ can be computationally problematic during recovery. Thus, rather than purely random Φ, we can also consider Φ that admit a fast transform-based implementation by taking random submatrices of a Walsh, Hadamard, or noiselet transform [103, 104]. We will describe the Walsh transform for the purpose of illustration.

We will suppose that N is a power of 2 and let $W_{\log_2 N}$ denote the $N \times N$ Walsh transform matrix. We begin by setting $W_0 = 1$, and we now define W_j recursively as

$$W_j = \frac{1}{\sqrt{2}} \begin{bmatrix} W_{j-1} & W_{j-1} \\ W_{j-1} & -W_{j-1} \end{bmatrix}.$$

This construction produces an orthonormal matrix with entries of $\pm 1/\sqrt{N}$ that admits a fast implementation requiring $O(N \log N)$ computations to apply. As an example, note that

$$W_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

and

$$W_2 = \frac{1}{2}\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}.$$

We can exploit these constructions as follows. Suppose that $N = 2^B$ and generate W_B. Let I_Γ denote a $M \times N$ submatrix of the identity I obtained by picking a random set of M rows, so that $I_\Gamma W_B$ is the submatrix of W_B consisting of the rows of W_B indexed by Γ. Furthermore, let D denote a random $N \times N$ permutation matrix. We can generate Φ as

$$\Phi = \left(\frac{1}{2}\sqrt{N} I_\Gamma W_B + \frac{1}{2}\right) D. \tag{5.1}$$

Note that $\frac{1}{2}\sqrt{N} I_\Gamma W_B + \frac{1}{2}$ merely rescales and shifts $I_\Gamma W_B$ to have $\{0, 1\}$-valued entries, and recall that each row of Φ will be reshaped into a 2-D matrix of numbers that is then displayed on the DMD array. Furthermore, D can be thought of as either permuting the pixels or permuting the columns of W_B. This step adds some additional randomness since some of the rows of the Walsh matrix are highly correlated with coarse scale wavelet basis functions — but permuting the pixels eliminates this structure. Note that at this point we do not have any strict guarantees that such Φ combined with a wavelet basis Ψ will yield a product $\Phi\Psi$ satisfying the RIP, but this approach seems to work well in practice.

5.2 The Random Demodulator

5.2.1 Architecture

We now turn to the question of acquiring compressive measurements of a continuous-time signal $x(t)$. Specifically, we would like to build an *analog-to-digital converter*

(ADC) that avoids having to sample $x(t)$ at its Nyquist rate when $x(t)$ is sparse. In this context, we will assume that $x(t)$ has some kind of sparse structure in the Fourier domain, meaning that it is still bandlimited but that much of the spectrum is empty. We will discuss the different possible signal models for mathematically capturing this structure in greater detail below. For now, the challenge is that our measurement system must be built using analog hardware. This imposes severe restrictions on the kinds of operations we can perform.

To be more concrete, since we are dealing with a continuous-time signal $x(t)$, we must also consider continuous-time test functions $\{\phi_j(t)\}_{j=1}^M$. We then consider a finite window of time, say $t \in [0, T]$, and would like to collect M measurements of the form

$$y[j] = \int_0^T x(t)\phi_j(t)\,dt. \tag{5.2}$$

Building an analog system to collect such measurements will require three main components:

1. hardware for generating the test signals $\phi_j(t)$;
2. M *correlators* that multiply the signal $x(t)$ with each respective $\phi_j(t)$; and
3. M *integrators* with a zero-valued initial state.

We could then sample and quantize the output of each of the integrators to collect the measurements $y[j]$. Of course, even in this somewhat idealized setting, it should be clear that what we can build in hardware will constrain our choice of $\phi_j(t)$ since we cannot reliably and accurately produce (and reproduce) arbitrarily complex $\phi_j(t)$ in analog hardware. Moreover, the architecture described above essentially requires M correlator/integrator pairs operating in parallel, which will be potentially prohibitively expensive both in terms of dollar cost as well as costs such as size, weight, and power (SWAP).

As a result, there have been a number of efforts to design simpler architectures, chiefly by carefully designing structured $\phi_j(t)$. The simplest to describe and historically earliest idea is to choose $\phi_j(t) = \delta(t - t_j)$, where $\{t_j\}_{j=1}^M$ denotes a sequence of M locations in time at which we would like to sample the signal $x(t)$. Typically, if the number of measurements we are acquiring is lower than the Nyquist-rate, then these locations cannot simply be uniformly spaced in the interval $[0, T]$, but must be carefully chosen. Note that this approach simply requires a single traditional ADC with the ability to sample on a non-uniform grid, avoiding the requirement for M parallel correlator/integrator pairs. Such non-uniform sampling systems have been studied in other contexts outside of the CS framework. For example, there exist specialized fast algorithms for the recovery of extremely large Fourier-sparse signals. The algorithm uses samples at a non-uniform sequence of locations that are highly structured, but where the initial location is chosen using a (pseudo)random seed. This literature provides guarantees similar to those available from standard CS [105, 106]. Additionally, there exist frameworks for the sampling and recovery of multi-band signals, whose Fourier transforms are mostly zero except for a few frequency bands. These schemes again use non-uniform sampling patterns based on coset sampling [23–27, 107]. Unfortunately, these approaches are often highly sensitive to *jitter*, or error in the timing of when the samples are taken.

We will consider a rather different approach, which we call the *random demodulator* [52, 108, 109].[2] The architecture of the random demodulator is depicted in Figure 5.4. The analog input $x(t)$ is correlated with a pseudorandom square pulse of ± 1s, called the *chipping sequence* $p_c(t)$, which alternates between values at a rate of N_aHz, where N_aHz is at least as fast as the Nyquist rate of $x(t)$. The mixed signal is integrated over a time period $1/M_a$ and sampled by a traditional integrate-and-dump

[2] A correlator is also known as a "demodulator" due to its most common application: demodulating radio signals.

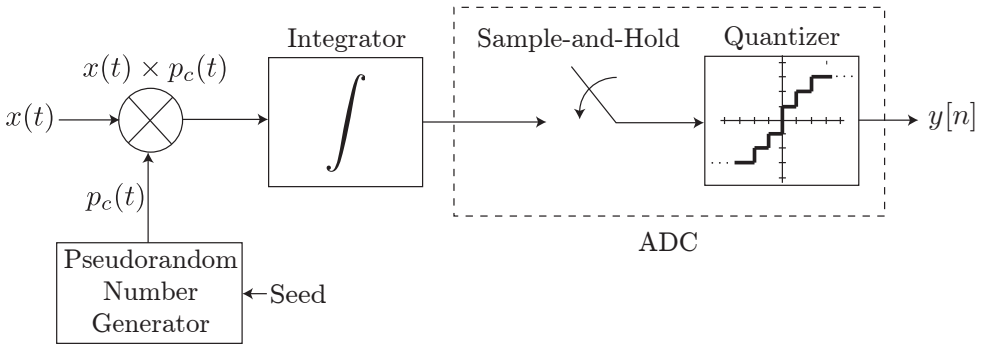

Figure 5.4: Random demodulator block diagram.

back-end ADC at M_aHz $\ll N_a$Hz. In this case our measurements are given by

$$y[j] = \int_{(j-1)/M_a}^{j/M_a} p_c(t)x(t)\,dt. \tag{5.3}$$

In practice, data is processed in time blocks of period T, and we define $N = N_aT$ as the number of elements in the chipping sequence, and $M = M_aT$ as the number of measurements. We will discuss the discretization of this model below, but the key observation is that the correlator and chipping sequence operate at a fast rate, while the back-end ADC operates at a low rate. In hardware it is easier to build a high-rate modulator/chipping sequence combination than a high-rate ADC [109]. In fact, many systems already use components of this front end for binary phase shift keying demodulation, as well as for other conventional communication schemes such as CDMA.

5.2.2 Discrete formulation

Although the random demodulator directly acquires compressive measurements without first sampling $x(t)$, it is equivalent to a system which first samples $x(t)$ at its Nyquist-rate to yield a discrete-time vector x, and then applies a matrix Φ to obtain the measurements $y = \Phi x$. To see this we let $p_c[n]$ denote the sequence of ± 1 used to

generate the signal $p_c(t)$, i.e., $p_c(t) = p_c[n]$ for $t \in [(n-1)/N_a, n/N_a]$. As an example, consider the first measurement, or the case of $j = 1$. In this case, $t \in [0, 1/M_a]$, so that $p_c(t)$ is determined by $p_c[n]$ for $n = 1, 2, \ldots, N_a/M_a$. Thus, from (5.3) we obtain

$$\begin{aligned} y[1] &= \int_0^{1/M_a} p_c(t)x(t)\,dt \\ &= \sum_{n=1}^{N_a/M_a} p_c[n] \int_{(n-1)/N_a}^{n/N_a} x(t)\,dt. \end{aligned}$$

But since N_a is the Nyquist-rate of $x(t)$, $\int_{(n-1)/N_a}^{n/N_a} x(t)\,dt$ simply calculates the average value of $x(t)$ on the n^{th} interval, yielding a sample denoted $x[n]$. Thus, we obtain

$$y[1] = \sum_{n=1}^{N_a/M_a} p_c[n]x[n].$$

In general, our measurement process is equivalent to multiplying the signal x with the random sequence of ± 1s in $p_c[n]$ and then summing every sequential block of N_a/M_a coefficients. We can represent this as a banded matrix Φ containing N_a/M_a pseudorandom ± 1s per row. For example, with $N = 12$, $M = 4$, and $T = 1$, such a Φ is expressed as

$$\Phi = \begin{bmatrix} -1 & +1 & +1 & & & & & & & & & \\ & & & -1 & +1 & -1 & & & & & & \\ & & & & & & +1 & +1 & -1 & & & \\ & & & & & & & & & +1 & -1 & -1 \end{bmatrix}. \tag{5.4}$$

In general, Φ will have M rows and each row will contain N/M nonzeros. Note that matrices satisfying this structure are extremely efficient to apply, requiring only $O(N)$ computations compared to $O(MN)$ in the general case. This is extremely useful during recovery.

A detailed analysis of the random demodulator in [52] studied the properties of

these matrices applied to a particular signal model. Specifically, it is shown that if Ψ represents the $N \times N$ normalized discrete Fourier transform (DFT) matrix, then the matrix $\Phi\Psi$ will satisfy the RIP with high probability, provided that

$$M = O\left(K \log^2(N/K)\right),$$

where the probability is taken with respect to the random choice of $p_c[n]$. This means that if $x(t)$ is a periodic (or finite-length) signal such that once it is sampled it is sparse or compressible in the basis Ψ, then it should be possible to recover $x(t)$ from the measurements provided by the random demodulator. Moreover, it is empirically demonstrated that combining ℓ_1 minimization with the random demodulator can recover K-sparse (in Ψ) signals with

$$M \geq CK \log(N/K + 1)$$

measurements where $C \approx 1.7$ [52].

Note that the signal model considered in [52] is somewhat restrictive, since even a pure tone will not yield a sparse DFT unless the frequency happens to be equal to k/N_a for some integer k. Perhaps a more realistic signal model is the multi-band signal model of [23–27, 107], where the signal is assumed to be bandlimited outside of K bands each of bandwidth B, where KB is much less than the total possible bandwidth. It remains unknown whether the random demodulator can be exploited to recover such signals. Moreover, there also exist other CS-inspired architectures that we have not explored in this section [53, 54, 110], and this remains an active area of research. We have simply provided an overview of one of the more promising approaches so as to illustrate the potential applicability of this thesis to the problem of analog-to-digital conversion.

Part III

Sparse Signal Recovery

Chapter 6

Sparse Recovery via Orthogonal Greedy Algorithms

We now turn to the problem of recovering sparse signals from the kind of measurements produced by the systems described in Part II. We begin by taking a closer look at some of the greedy algorithms described in Sections 2.4.2 and 2.4.3 in the context of matrices satisfying the RIP. Specifically, in this chapter[1] we provide an RIP-based theoretical framework for analyzing orthogonal greedy algorithms. First, we provide an RIP-based analysis of the classical algorithm of OMP when applied to recovering sparse signals in the noise-free setting. This analysis revolves around three key facts: *(i)* that in each step of the algorithm, the residual vector r^ℓ can be written as a matrix times a sparse signal, *(ii)* that this matrix satisfies the RIP, and *(iii)* that consequently a sharp bound can be established for the vector h^ℓ of inner products. Our main conclusion, Theorem 6.1, states that the RIP of order $K+1$ (with $\delta < 1/(3\sqrt{K})$) is sufficient for OMP to exactly recover any K-sparse signal in exactly K iterations. However, for restricted classes of K-sparse signals (those with sufficiently strong decay in the nonzero coefficients), a relaxed bound on the isometry

[1]This work was done in collaboration with Michael B. Wakin [111].

constant can be used. We then extend this analysis and use the same techniques to establish a simple proof that under even weaker assumptions, ROMP will also succeed in recovering K-sparse signals.

6.1 Orthogonal Matching Pursuit

Theoretical analysis of OMP to date has concentrated primarily on two fronts. The first has involved the notion of a coherence parameter $\mu := \max_{i,j} |\langle \phi_i, \phi_j \rangle|$, where ϕ_i denotes column i of the matrix Φ. When the columns of Φ have unit norm and $\mu < 1/(2K-1)$, it has been shown [76] that OMP will recover any $x \in \Sigma_K$ from the (noise-free) measurements $y = \Phi x$. This guarantee is deterministic and applies to any matrix Φ having normalized columns and $\mu < 1/(2K-1)$.

The second analytical front has involved the notion of probability. Suppose that $x \in \Sigma_K$ and that Φ is drawn from a suitable random distribution (independently of x) with $M = O(K \log(N))$ rows. Then with high probability, OMP will recover x exactly from the measurements $y = \Phi x$ [45]. It is not guaranteed, however, that any such fixed matrix will allow recovery of all sparse x simultaneously.

As an alternative to coherence and to probabilistic analysis, a large number of algorithms within the broader field of sparse recovery have been studied using the RIP as described in Chapter 3. As noted in Section 3.3, when it is satisfied, the RIP for a matrix Φ provides a sufficient condition to guarantee successful sparse recovery using a wide variety of algorithms [30, 32, 33, 35, 42–44, 77, 78, 112]. Nevertheless, despite the considerable attention that has been paid to both OMP and the RIP, analysis of OMP using the RIP has been relatively elusive to date. However, several alternative greedy algorithms have been proposed — all essentially modifications of OMP — that are apparently much more amenable to RIP-based analysis. Specifically, both ROMP and CoSaMP, as well as Subspace Pursuit (SP) [35] and DThresh [33], are essentially

all extensions of OMP that have been tweaked primarily to enable their analysis using the RIP. For each of these algorithms it has been shown that the RIP of order CK (where $C \geq 2$ is a constant depending on the algorithm) with δ adequately small is sufficient for exact recovery of K sparse signals. In this chapter we show that the original formulation of OMP also satisfies this property.

Towards this end, we begin with some very simple observations regarding OMP as presented in Algorithm 3. The key idea is to try to iteratively estimate a set Λ that contains the locations of the nonzeros of x by starting with $\Lambda = \emptyset$ and then adding a new element to Λ in each iteration. In order to select which element to add, the algorithm also maintains a residual vector $r \notin \mathcal{R}(\Phi_\Lambda)$ that represents the component of the measurement vector y that cannot be explained by the columns of Φ_Λ. Specifically, at the beginning of the ℓ^{th} iteration, Λ^ℓ is our current estimate of $\text{supp}(x)$, and the residual r^ℓ is defined as $r^\ell = y - \Phi x^\ell$ where $\text{supp}(x^\ell) \subseteq \Lambda^\ell$. The element added to Λ^ℓ is the index of the column of Φ that has the largest inner product with r^ℓ.

Our first observation is that r^ℓ can be viewed as the orthogonalization of y against the previously chosen columns of Φ. To see this, note that the solution to the least squares problem in the update step is given by

$$\begin{aligned} x^\ell|_{\Lambda^\ell} &= \Phi_{\Lambda^\ell}^\dagger y \\ x^\ell|_{(\Lambda^\ell)^c} &= 0. \end{aligned} \tag{6.1}$$

Thus we observe that

$$r^\ell = y - \Phi x^\ell = y - \Phi_{\Lambda^\ell}\Phi_{\Lambda^\ell}^\dagger y = (I - P_{\Lambda^\ell})y = P_{\Lambda^\ell}^\perp y.$$

Note that it is not actually necessary to explicitly compute x^ℓ in order to calculate r^ℓ.

For our second observation, we define $A_\Lambda := P_\Lambda^\perp \Phi$. This matrix is the result of orthogonalizing the columns of Φ against $\mathcal{R}(\Phi_\Lambda)$. It is therefore equal to zero on columns indexed by Λ. Note that in the proxy step, one may correlate r^ℓ either with the columns of Φ or with the columns of A_{Λ^ℓ}. To see this equivalence, observe that $r^\ell = P_{\Lambda^\ell}^\perp y = P_{\Lambda^\ell}^\perp P_{\Lambda^\ell}^\perp y = (P_{\Lambda^\ell}^\perp)^T P_{\Lambda^\ell}^\perp y$ and so

$$h^\ell = \Phi^T r^\ell = \Phi^T (P_{\Lambda^\ell}^\perp)^T P_{\Lambda^\ell}^\perp y = A_{\Lambda^\ell}^T r^\ell. \tag{6.2}$$

Incidentally, along these same lines we observe that

$$h^\ell = \Phi^T r^\ell = \Phi^T P_{\Lambda^\ell}^\perp y = \Phi^T (P_{\Lambda^\ell}^\perp)^T y = A_{\Lambda^\ell}^T y.$$

From this we note that it is not actually necessary to explicitly compute r^ℓ in order to calculate the inner products during the proxy step; in fact, the original formulation of OMP was stated with instructions to orthogonalize the remaining columns of Φ against those previously chosen and merely correlate the resulting vectors against y [71, 75]. Additionally, we recall that, in A_{Λ^ℓ}, all columns indexed by Λ^ℓ will be zero. It follows that

$$h^\ell(j) = 0 \quad \forall j \in \Lambda^\ell, \tag{6.3}$$

and so, since $\Lambda^\ell = \Lambda^{\ell-1} \cup \{j^*\}$ with $j^* \notin \Lambda^{\ell-1}$,

$$|\Lambda^\ell| = \ell. \tag{6.4}$$

Our third observation is that, in the case of noise-free measurements $y = \Phi x$, we may write

$$r^\ell = P_{\Lambda^\ell}^\perp y = P_{\Lambda^\ell}^\perp \Phi x = A_{\Lambda^\ell} x.$$

Again recalling that all columns of A_{Λ^ℓ} indexed by Λ^ℓ are zero, we thus note that

when $\text{supp}(x) \subseteq \Lambda^\ell$, $r^\ell = 0$, and from (6.1) we also know that $x^\ell = x$ exactly. It will also be useful to note that for the same reason, we can also write

$$r^\ell = A_{\Lambda^\ell} \widetilde{x}^\ell, \tag{6.5}$$

where

$$\widetilde{x}^\ell|_{\Lambda^\ell} = 0 \quad \text{and} \quad \widetilde{x}^\ell|_{(\Lambda^\ell)^c} = x|_{(\Lambda^\ell)^c}. \tag{6.6}$$

6.2 Analysis of OMP

Our analysis of OMP will center on the vector h^ℓ. In light of (6.2) and (6.5), we see that A_{Λ^ℓ} plays a role both in *constructing* and in *analyzing* the residual vector. In Lemma 6.2 below, we show that the matrix A_{Λ^ℓ} satisfies a modified version of the RIP. This allows us to very precisely bound the values of the inner products in the vector h^ℓ.

We begin with two elementary lemmas which will have repeated applications throughout this thesis. Our first result, which is a straightforward generalization of Lemma 2.1 of [78], states that RIP operators approximately preserve inner products between sparse vectors.

Lemma 6.1. *Let $u, v \in \mathbb{R}^N$ be given, and suppose that a matrix Φ satisfies the RIP of order $\max(\|u+v\|_0, \|u-v\|_0)$ with isometry constant δ. Then*

$$|\langle \Phi u, \Phi v \rangle - \langle u, v \rangle| \leq \delta \|u\|_2 \|v\|_2. \tag{6.7}$$

Proof. We first assume that $\|u\|_2 = \|v\|_2 = 1$. From the fact that

$$\|u \pm v\|_2^2 = \|u\|_2^2 + \|v\|_2^2 \pm 2\langle u, v \rangle = 2 \pm 2\langle u, v \rangle$$

and since Φ satisfies the RIP, we have that

$$(1-\delta)(2 \pm 2\langle u, v\rangle) \leq \|\Phi u \pm \Phi v\|_2^2 \leq (1+\delta)(2 \pm 2\langle u, v\rangle).$$

From the parallelogram identity we obtain

$$\begin{aligned}\langle \Phi u, \Phi v\rangle &= \frac{1}{4}\left(\|\Phi u + \Phi v\|_2^2 - \|\Phi u - \Phi v\|_2^2\right) \\ &\leq \frac{(1+\langle u, v\rangle)(1+\delta) - (1-\langle u, v\rangle)(1-\delta)}{2} \\ &= \langle u, v\rangle + \delta.\end{aligned}$$

Similarly, one can show that $\langle \Phi u, \Phi v\rangle \geq \langle u, v\rangle - \delta$, and thus $|\langle \Phi u, \Phi v\rangle - \langle u, v\rangle| \leq \delta$. The result follows for u, v with arbitrary norm from the bilinearity of the inner product. □

One consequence of this result is that sparse vectors that are orthogonal in $\mathbb{R}^N$ remain nearly orthogonal after the application of Φ. From this observation, it can be demonstrated[2] that if Φ has the RIP, then A_Λ satisfies a modified version of the RIP.

Lemma 6.2. *Suppose that Φ satisfies the RIP of order K with isometry constant δ, and let $\Lambda \subset \{1, 2, \ldots, N\}$. If $|\Lambda| < K$ then*

$$\left(1 - \frac{\delta}{1-\delta}\right)\|u\|_2^2 \leq \|A_\Lambda u\|_2^2 \leq (1+\delta)\|u\|_2^2 \tag{6.8}$$

for all $u \in \mathbb{R}^N$ such that $\|u\|_0 \leq K - |\Lambda|$ and $\operatorname{supp}(u) \cap \Lambda = \emptyset$.

Proof. From the definition of A_Λ we may decompose $A_\Lambda u$ as $A_\Lambda u = \Phi u - P_\Lambda \Phi u$. Since P_Λ is an orthogonal projection, we can write

$$\|\Phi u\|_2^2 = \|P_\Lambda \Phi u\|_2^2 + \|A_\Lambda u\|_2^2. \tag{6.9}$$

[2]This result was first proven in [113] and then independently in [35].

Our goal is to show that $\|\Phi u\|_2 \approx \|A_\Lambda u\|_2$, or equivalently, that $\|P_\Lambda \Phi u\|_2$ is small. Towards this end, we note that since $P_\Lambda \Phi u$ is orthogonal to $A_\Lambda u$,

$$\begin{aligned}\langle P_\Lambda \Phi u, \Phi u\rangle &= \langle P_\Lambda \Phi u, P_\Lambda \Phi u + A_\Lambda u\rangle \\ &= \langle P_\Lambda \Phi u, P_\Lambda \Phi u\rangle + \langle P_\Lambda \Phi u, A_\Lambda u\rangle \\ &= \|P_\Lambda \Phi u\|_2^2. \end{aligned} \tag{6.10}$$

Since P_Λ is a projection onto $\mathcal{R}(\Phi_\Lambda)$ there exists a $z \in \mathbb{R}^N$ with $\mathrm{supp}(z) \subseteq \Lambda$ such that $P_\Lambda \Phi u = \Phi z$. Furthermore, by assumption, $\mathrm{supp}(u) \cap \Lambda = \emptyset$. Hence $\langle u, z\rangle = 0$ and from the RIP and Lemma 6.1,

$$\frac{|\langle P_\Lambda \Phi u, \Phi u\rangle|}{\|P_\Lambda \Phi u\|_2 \|\Phi u\|_2} = \frac{|\langle \Phi z, \Phi u\rangle|}{\|\Phi z\|_2 \|\Phi u\|_2} \leq \frac{|\langle \Phi z, \Phi u\rangle|}{(1-\delta)\|z\|_2 \|u\|_2} \leq \frac{\delta}{1-\delta}.$$

Combining this with (6.10), we obtain

$$\|P_\Lambda \Phi u\|_2 \leq \frac{\delta}{1-\delta} \|\Phi u\|_2.$$

Since we trivially have that $\|P_\Lambda \Phi u\|_2 \geq 0$, we can combine this with (6.9) to obtain

$$\left(1 - \left(\frac{\delta}{1-\delta}\right)^2\right) \|\Phi u\|_2^2 \leq \|A_\Lambda u\|_2^2 \leq \|\Phi u\|_2^2.$$

Since $\|u\|_0 \leq K$, we can use the RIP to obtain

$$\left(1 - \left(\frac{\delta}{1-\delta}\right)^2\right) (1-\delta)\|u\|_2^2 \leq \|A_\Lambda u\|_2^2 \leq (1+\delta)\|u\|_2^2,$$

which simplifies to (6.8). □

In other words, if Φ satisfies the RIP of order K, then A_Λ acts as an approximate isometry on every $(K - |\Lambda|)$-sparse vector supported on Λ^c. From (6.5), we recall that

the residual vector in OMP is formed by applying A_{Λ^ℓ} to a sparse vector supported on $(\Lambda^\ell)^c$. Combining the above results, then, we may bound the inner products $h^\ell(j)$ as follows.

Lemma 6.3. *Let* $\Lambda \subset \{1, 2, \ldots, N\}$ *and suppose* $\widetilde{x} \in \mathbb{R}^N$ *with* $\mathrm{supp}(\widetilde{x}) \cap \Lambda = \emptyset$. *Define*

$$h = A_\Lambda^T A_\Lambda \widetilde{x}. \tag{6.11}$$

Then if Φ *satisfies the RIP of order* $\|\widetilde{x}\|_0 + |\Lambda| + 1$ *with isometry constant* δ, *we have*

$$|h(j) - \widetilde{x}(j)| \leq \frac{\delta}{1-\delta} \|\widetilde{x}\|_2 \tag{6.12}$$

for all $j \notin \Lambda$.

Proof. From Lemma 6.2 we have that the restriction of A_Λ to the columns indexed by Λ^c satisfies the RIP of order $(\|\widetilde{x}\|_0 + |\Lambda| + 1) - |\Lambda| = \|\widetilde{x}\|_0 + 1$ with isometry constant $\delta/(1-\delta)$. By the definition of h, we also know that

$$h(j) = \langle A_\Lambda \widetilde{x}, A_\Lambda e_j \rangle,$$

where e_j denotes the j^{th} vector from the cardinal basis. Now, suppose $j \notin \Lambda$. Then because $\|\widetilde{x} \pm e_j\|_0 \leq \|\widetilde{x}\|_0 + 1$ and $\mathrm{supp}(\widetilde{x} \pm e_j) \cap \Lambda = \emptyset$, we conclude from Lemma 6.1 that

$$|h(j) - \widetilde{x}(j)| = |\langle A_\Lambda \widetilde{x}, A_\Lambda e_j \rangle - \langle \widetilde{x}, e_j \rangle| \leq \frac{\delta}{1-\delta} \|\widetilde{x}\|_2 \|e_j\|_2.$$

Noting that $\|e_j\|_2 = 1$, we reach the desired conclusion. □

With this bound on the inner products $h^\ell(j)$, we may derive a sufficient condition under which the identification step of OMP will succeed.

Corollary 6.1. *Suppose that Λ, Φ, $\widetilde{x}$ meet the assumptions specified in Lemma 6.3, and let h be as defined in (6.11). If*

$$\|\widetilde{x}\|_\infty > \frac{2\delta}{1-\delta}\|\widetilde{x}\|_2, \tag{6.13}$$

we are guaranteed that $\arg\max_j |h(j)| \in \mathrm{supp}(\widetilde{x})$.

Proof. If (6.12) is satisfied, then for indices $j \notin \mathrm{supp}(\widetilde{x})$, we will have

$$|h(j)| \leq \frac{\delta}{1-\delta}\|\widetilde{x}\|_2.$$

(Recall from (6.3) that $h(j) = 0$ for $j \in \Lambda$.) If (6.13) is satisfied, then there exists some $j \in \mathrm{supp}(\widetilde{x})$ with

$$|\widetilde{x}(j)| > \frac{2\delta}{1-\delta}\|\widetilde{x}\|_2.$$

From (6.12) and the triangle inequality, we conclude that for this index j,

$$|h(j)| > \frac{\delta}{1-\delta}\|\widetilde{x}\|_2.$$

Thus, we have that $\max_j |h(j)| > |h(k)|$ for all $k \notin \mathrm{supp}(\widetilde{x})$, which ensures that $\arg\max_j |h(j)| \in \mathrm{supp}(\widetilde{x})$, as desired. □

By choosing δ small enough, it is possible to guarantee that the condition (6.13) is satisfied. In particular, the lemma below follows from standard arguments.

Lemma 6.4. *For any $u \in \mathbb{R}^N$, $\|u\|_\infty \geq \|u\|_2/\sqrt{\|u\|_0}$.*

Proof. Since we can bound $|u_j| \leq \|u\|_\infty$ for all j, we have that

$$\|u\|_2 = \sqrt{\sum_{j\in\mathrm{supp}(u)} |u_j|^2} \leq \sqrt{\sum_{j\in\mathrm{supp}(u)} \|u\|_\infty^2} = \sqrt{\|u\|_0}\|u\|_\infty,$$

as desired. □

Putting these results together, we can now establish our main theorem concerning OMP.

Theorem 6.1. *Suppose that Φ satisfies the RIP of order $K+1$ with isometry constant $\delta < 1/(3\sqrt{K})$. Then for any $x \in \Sigma_K$, OMP will recover x exactly from $y = \Phi x$ in K iterations.*

Proof. The proof works by induction. We start with the first iteration where $h^0 = \Phi^T \Phi x$ and note that $\Phi = A_\emptyset$. Because $\|x\|_0 \leq K$, Lemma 6.4 states that $\|x\|_\infty \geq \|x\|_2/\sqrt{K}$. One can also check that $\delta < 1/(3\sqrt{K})$ implies that

$$\frac{2\delta}{1-\delta} < \frac{1}{\sqrt{K}}.$$

Therefore, we are guaranteed that (6.13) is satisfied, and so from Corollary 6.1 we conclude that $\arg\max_j |h^0(j)| \in \text{supp}(x)$.

We now consider the general induction step. Suppose that we are at iteration ℓ and that all previous iterations have succeeded, by which we mean that $\Lambda^\ell \subseteq \text{supp}(x)$. From (6.6), we know that $\text{supp}(\widetilde{x}^\ell) \cap \Lambda^\ell = \emptyset$ and that $\|\widetilde{x}^\ell\|_0 \leq K - \ell$. From (6.4), we know that $|\Lambda^\ell| = \ell$. By assumption, Φ satisfies the RIP of order $K + 1 = (K - \ell) + \ell + 1 \geq \|\widetilde{x}^\ell\|_0 + |\Lambda^\ell| + 1$. Finally, using Lemma 6.4, we have that

$$\|\widetilde{x}^\ell\|_\infty \geq \frac{\|\widetilde{x}^\ell\|_2}{\sqrt{K-\ell}} \geq \frac{\|\widetilde{x}^\ell\|_2}{\sqrt{K}} > \frac{2\delta}{1-\delta}\|\widetilde{x}^\ell\|_2.$$

From Corollary 6.1 we conclude that $\arg\max_j |h^\ell(j)| \in \text{supp}(\widetilde{x}^\ell)$ and hence $\Lambda^{\ell+1} \subseteq \text{supp}(x)$. □

6.3 Context

Let us place Theorem 6.1 in the context of the OMP literature. Using the RIP as a sufficient condition to guarantee OMP performance is apparently novel. Moreover,

the fact that our bound requires only the RIP of order $K+1$ is apparently unique among the published CS literature; much more common are results requiring the RIP of order $1.75K$ [112], $2K$ [42, 78], $3K$ [30, 35], $4K$ [44], and so on. Of course, such results often permit the isometry constant to be much larger.[3]

If one wishes to use the RIP of order $K+1$ as a sufficient condition for exact recovery of all K-sparse signals via OMP (as we have), then little improvement is possible in relaxing the isometry constant δ above $1/(3\sqrt{K})$. In particular, there exists a matrix satisfying the RIP of order $K+1$ with $\delta \leq 1/\sqrt{K}$ for which there exists a K-sparse signal $x \in \mathbb{R}^N$ that cannot be recovered exactly via K iterations of OMP. (This is conjectured in [35] with a suggestion for constructing such a matrix, and for the case $K=2$ we have confirmed this via experimentation.)

Unfortunately, Theorem 4.3 suggests that finding a matrix Φ satisfying the RIP of order $K+1$ with an isometry constant $\delta < 1/(3\sqrt{K})$ will possibly require $M = O(K^2 \log(N/K))$ random measurements. In fact, Theorem 3.4 tells us that this RIP condition necessitates that we at least have $M = O(K^{3/2})$. However, if one wishes to guarantee exact recovery of all K-sparse signals via OMP (as we have), then there is little room for further reducing M. In particular, it has been argued in a recent paper concerned with uniform guarantees for greedy algorithms [115] that there exists a constant C such that when $M \leq CK^{3/2}$, for most random $M \times N$ matrices Φ there will exist some K-sparse signal $x \in \mathbb{R}^N$ that cannot be recovered exactly via K iterations of OMP.

It is also worth comparing our RIP-based analysis with coherence-based analysis [76], as both techniques provide a sufficient condition for OMP to recover all K-sparse signals. It has been shown [45] that in a random $M \times N$ matrix, the co-

[3]Recently, it was shown in [114] that the RIP of order K with $\delta < 0.307$ is a sufficient condition for recovery via ℓ_1 minimization in the absence of noise. In general it is important to note that a smaller order of the RIP is not necessarily a weaker requirement if the required constant is also significantly smaller due to Lemma 3.1. For example, Lemma 3.1 implies that if Φ satisfies the RIP of order $K+1$ with constant δ, then Φ also satisfies the RIP of order $2K$ with constant 4δ.

herence parameter μ is unlikely to be smaller than $\log(N)/\sqrt{M}$. Thus, to ensure $\mu < 1/(2K-1)$, one requires $M = O(K^2 \log^2(N))$, which is roughly the same as what is required by our analysis. We note that neither result is strictly stronger than the other; we have confirmed experimentally that there exist matrices that satisfy our RIP condition but not the coherence condition, and vice versa.

Finally, we note that the aforementioned modifications of OMP (the ROMP, SP, CoSaMP, and DThresh algorithms) all have RIP-based guarantees of robust recovery in noise and stable recovery of non-sparse signals. Until recently, no such RIP-based or coherence-based guarantees had been established for OMP itself. However, there has been recent progress in using the RIP and similar conditions to analyze the performance of OMP on non-sparse signals [116]. The results of [116] can be adapted to provide a guarantee of exact recovery for sparse signals, but the assumptions required are stronger than the assumptions made in this work. Furthermore, a number of additional open questions remain concerning the performance of OMP on non-sparse signals, and performance in the presence of noise has yet to be fully addressed. We speculate that our perspective may help to further the general understanding of OMP and perhaps provide a route to such guarantees. At present, however, this remains a topic of ongoing work [116–119].

6.4 Extensions

6.4.1 Strongly-decaying sparse signals

For even moderate values of the isometry constant δ there exist sparse signals that we can ensure are recovered exactly. For example, if the decay of coefficients is sufficiently strong in a sparse signal, we may use Lemma 6.3 to ensure that the signal entries are recovered in the order of their magnitude.

For any $x \in \mathbb{R}^N$ with $\|x\|_0 \leq K$ we denote by $x'(j)$ the entries of x ordered by

magnitude, i.e.,

$$|x'(1)| \geq |x'(2)| \geq \cdots \geq |x'(K)| \geq 0$$

with $x'(K+1) = x'(K+2) = \cdots = x'(N) = 0$.

Theorem 6.2. *Suppose that Φ satisfies the RIP of order $K+1$ with isometry constant $\delta < 1/3$. Suppose $x \in \Sigma_K$ and that for all $j \in \{1, 2, \ldots, K-1\}$,*

$$\frac{|x'(j)|}{|x'(j+1)|} \geq \alpha.$$

If

$$\alpha > \frac{1 + 2\frac{\delta}{1-\delta}\sqrt{K-1}}{1 - 2\frac{\delta}{1-\delta}}, \tag{6.14}$$

then OMP will recover x exactly from $y = \Phi x$ in K iterations.

Proof. The proof again proceeds by induction. At each stage, OMP will choose the largest entry of $\widetilde{x}^\ell$. To see this, note that by (6.12) we have

$$|h^\ell(j) - \widetilde{x}^\ell(j)| \leq \frac{\delta}{1-\delta}\|\widetilde{x}^\ell\|_2.$$

The nonzero entries of $\widetilde{x}^\ell$ will be comprised of $x'(\ell+1), x'(\ell+2), \ldots, x'(K)$. Thus,

$$\begin{aligned}
\|\widetilde{x}^\ell\|_2 &\leq \sqrt{|x'(\ell+1)|^2 + (K-1)\frac{|x'(\ell+1)|^2}{\alpha^2}} \\
&= \frac{|x'(\ell+1)|}{\alpha}\sqrt{\alpha^2 + (K-1)} \\
&\leq \frac{|x'(\ell+1)|}{\alpha}(\alpha + \sqrt{K-1}).
\end{aligned}$$

Now, for the specific index j at which $\widetilde{x}^\ell$ has its largest entry, we have

$$\begin{aligned}
|h^\ell(j)| &\geq |x'(\ell+1)| - \frac{\delta}{1-\delta}\frac{|x'(\ell+1)|}{\alpha}\left(\alpha + \sqrt{K-1}\right) \\
&= \frac{|x'(\ell+1)|}{\alpha}\left(\alpha - \frac{\delta}{1-\delta}\left(\alpha + \sqrt{K-1}\right)\right),
\end{aligned} \tag{6.15}$$

while for all other values of j we have

$$\begin{aligned} |h^\ell(j)| &\leq |x'(\ell+2)| + \frac{\delta}{1-\delta}\frac{|x'(\ell+1)|}{\alpha}\left(\alpha + \sqrt{K-1}\right) \\ &\leq \frac{|x'(\ell+1)|}{\alpha}\left(1 + \frac{\delta}{1-\delta}\left(\alpha + \sqrt{K-1}\right)\right). \end{aligned} \tag{6.16}$$

From (6.14), it follows that (6.15) is greater than (6.16). □

6.4.2 Analysis of other orthogonal greedy algorithms

We now demonstrate that the techniques used above can also be used to analyze other orthogonal greedy algorithms. We focus on ROMP for the purpose of illustration, but similar methods should be able to simplify the analysis of other orthogonal greedy algorithms such as SP.[4]

We first briefly recall the difference between ROMP and OMP, which lies only in the identification step: whereas OMP adds only one index to Λ^ℓ at each iteration, ROMP adds up to K indices to Λ^ℓ at each iteration. Specifically, ROMP first selects the indices corresponding to the K largest elements in magnitude of h^ℓ (or all nonzero elements of h^ℓ if h^ℓ has fewer than K nonzeros), and denotes this set as Ω^ℓ. The next step is to regularize this set so that the values are comparable in magnitude. To do this, we define $R(\Omega^\ell) := \{\Omega \subseteq \Omega^\ell : |h^\ell(i)| \leq 2|h^\ell(j)| \ \forall i,j \in \Omega\}$, and set

$$\Omega_0^\ell := \arg\max_{\Omega \in R(\Omega^\ell)} \|h^\ell|_\Omega\|_2,$$

i.e., Ω_0^ℓ is the set with maximal energy among all regularized subsets of Ω^ℓ. Finally, setting $\Lambda^{\ell+1} = \Lambda^\ell \cup \Omega_0^\ell$, the remainder of the ROMP algorithm is identical to OMP.

[4]Some of the greedy algorithms that have been proposed recently, such as CoSaMP and DThresh, do not orthogonalize the residual against the previously chosen columns at each iteration, and so the techniques above cannot be directly applied to these algorithms. However, this orthogonalization step could easily be added (which in the case of CoSaMP yields an algorithm nearly identical to SP). Orthogonalized versions of these algorithms could then be studied using these techniques.

In order to analyze ROMP, we will need only two preliminary lemmas from [42], which we state without proof. Note that Lemma 6.5, which is essentially a generalization of Lemma 6.3, is stated using slightly weaker assumptions than are used in [42] and, to be consistent with the rest of this thesis, uses the quadratic form of the RIP (whereas [42] uses the non-quadratic form). However, the present version can easily be obtained using the same proof techniques.

Lemma 6.5 (Needell-Vershynin [42])**.** *Let $\Gamma \subset \{1, 2, \ldots, N\}$ and $x \in \mathbb{R}^N$ be given. Then if Φ satisfies the RIP of order $|\mathrm{supp}(x) \cup \Gamma|$ with isometry constant δ, we have*

$$\|(\Phi^T \Phi x)|_\Gamma - x|_\Gamma\|_2 \leq \delta \|x\|_2.$$

Lemma 6.6 (Needell-Vershynin [42])**.** *Let $u \in \mathbb{R}^K$, $K > 1$, be arbitrary. Then there exists a subset $\Gamma \subseteq \{1, \ldots, K\}$ such that $|u(i)| \leq 2|u(j)|$ for all $i, j \in \Gamma$ and*

$$\|u|_\Gamma\|_2 \geq \frac{1}{2.5\sqrt{\log_2 K}} \|u\|_2.$$

Using these lemmas, we now provide a simplified proof of the main result of [42] concerning the recovery of sparse signals using ROMP.[5]

Theorem 6.3. *Suppose that Φ satisfies the RIP of order $3K$ with isometry constant $\delta \leq 0.13/\sqrt{\log_2 K}$. Then for any $x \in \Sigma_K$, ROMP will recover x exactly from $y = \Phi x$ in at most K iterations.*

Proof. The proof works by showing that at each iteration,

$$|\Omega_0^\ell \cap \mathrm{supp}(x)| \geq \frac{1}{2} |\Omega_0^\ell|. \tag{6.17}$$

[5]Note that we assume that Φ satisfies the RIP of order $3K$ with constant $\delta \leq 0.13/\sqrt{\log_2 K}$. Using Lemma 3.1, we can replace this with the assumption that Φ satisfies the RIP of order $2K$ with constant $\delta \leq .043/\sqrt{\log_2 K}$.

If (6.17) is satisfied for $0, 1, \ldots, \ell - 1$, then at iteration ℓ we have that

$$|\Lambda^\ell \cap \operatorname{supp}(x)| \geq \frac{1}{2}|\Lambda^\ell|. \tag{6.18}$$

It follows that, before $|\Lambda^\ell|$ exceeds $2K$, we will have $\operatorname{supp}(x) \subseteq \Lambda^\ell$. Because Φ satisfies the RIP of order $3K > 2K$, at termination, Φ_{Λ^ℓ} will be full rank. From (6.1) we conclude that $x^\ell = x$ exactly.

To prove (6.17), we again proceed by induction. Hence, we assume that (6.17) holds for $0, 1, \ldots, \ell - 1$, and thus (6.18) holds for iteration ℓ. We next assume for the sake of a contradiction that (6.17) does not hold for iteration ℓ, i.e., that

$$|\Omega_0^\ell \setminus \operatorname{supp}(x)| > \frac{1}{2}|\Omega_0^\ell|. \tag{6.19}$$

Define the sets $T = \Omega_0^\ell \setminus \operatorname{supp}(x)$ and $S = \operatorname{supp}(x) \setminus \Lambda^\ell = \operatorname{supp}(\widetilde{x}^\ell)$, where $\widetilde{x}^\ell$ is defined as in (6.6). Recall that we can write $h^\ell = A_{\Lambda^\ell}^T A_{\Lambda^\ell} \widetilde{x}^\ell$. Thus, using the assumption that $|T| > \frac{1}{2}|\Omega_0^\ell|$ and the facts that $T \subseteq \Omega_0^\ell$ and $\Omega_0^\ell \in R(\Omega^\ell)$, one can show that

$$\|h^\ell|_T\|_2 \geq \frac{1}{\sqrt{5}}\|h^\ell|_{\Omega_0^\ell}\|_2. \tag{6.20}$$

We now observe that

$$\|h^\ell|_{\Omega_0^\ell}\|_2 \geq \frac{1}{2.5\sqrt{\log_2 K}}\|h^\ell|_{\Omega^\ell}\|_2, \tag{6.21}$$

which follows from Lemma 6.6 and the fact that Ω_0^ℓ is the maximal regularizing set. From the maximality of Ω^ℓ and the fact that $|S| \leq K$, we have that $\|h^\ell|_{\Omega^\ell}\|_2 \geq \|h^\ell|_S\|_2$, so that by combining (6.20) and (6.21) we obtain

$$\|h^\ell|_T\|_2 \geq \frac{1}{2.5\sqrt{5\log_2 K}}\|h^\ell|_S\|_2. \tag{6.22}$$

Note that $|S \cup \operatorname{supp}(\widetilde{x}^\ell)| = |S| \leq K$ and since $|\Lambda^\ell| \leq 2K$, from Lemma 6.2 we have

that A_{Λ^ℓ} satisfies the RIP of order at least K with constant $\delta/(1-\delta)$, thus Lemma 6.5 implies that

$$\|h^\ell|_S - \widetilde{x}^\ell|_S\|_2 \leq \frac{\delta}{1-\delta}\|\widetilde{x}^\ell\|_2.$$

Since $\widetilde{x}^\ell|_{S^c} = 0$, $\|h^\ell|_S - \widetilde{x}^\ell|_S\|_2 \geq \|\widetilde{x}^\ell\|_2 - \|h^\ell|_S\|_2$, and thus

$$\|h^\ell|_S\|_2 \geq \frac{1-2\delta}{1-\delta}\|\widetilde{x}^\ell\|_2.$$

Hence,

$$\|h^\ell|_T\|_2 \geq \frac{(1-2\delta)/(1-\delta)}{2.5\sqrt{5\log_2 K}}\|\widetilde{x}^\ell\|_2. \tag{6.23}$$

On the other hand, since $|\mathrm{supp}(\widetilde{x}^\ell)| + |\Lambda^\ell \cap \mathrm{supp}(x)| = K$, from (6.18) we obtain that $|\mathrm{supp}(\widetilde{x}^\ell)| \leq K - |\Lambda^\ell|/2$. Thus, $|T \cup \mathrm{supp}(\widetilde{x}^\ell)| \leq |T| + |\mathrm{supp}(\widetilde{x}^\ell)| \leq 2K - |\Lambda^\ell|/2$. Furthermore, A_{Λ^ℓ} satisfies the RIP of order $3K - |\Lambda^\ell| = 3K - |\Lambda^\ell|/2 - |\Lambda^\ell|/2$. Since $|\Lambda^\ell| \leq 2K$, we have that A_{Λ^ℓ} satisfies the RIP of order at least $2K - |\Lambda^\ell|/2$ with constant $\delta/(1-\delta)$. Thus, Lemma 6.5 also implies that

$$\|h^\ell|_T\|_2 = \|h^\ell|_T - \widetilde{x}^\ell|_T\|_2 \leq \frac{\delta}{1-\delta}\|\widetilde{x}^\ell\|_2. \tag{6.24}$$

This is a contradiction whenever the right-hand-side of (6.23) is greater than the right-hand-side of (6.24), which occurs when $\delta < 1/(2 + 2.5\sqrt{5\log_2 K})$. Since $\log_2 K \geq 1$, we can replace this with the slightly stricter condition

$$\delta < 1/((2+2.5\sqrt{5})\sqrt{\log_2 K}) \approx 0.1317/\sqrt{\log_2 K}.$$

□

Observe that when $K = 1$, this proof (as well as the proofs in [42, 43]) break down since Lemma 6.6 does not apply. However, when $K = 1$ the ROMP algorithm simply reduces to OMP. In this case we can apply Theorem 6.1 to verify that ROMP

succeeds when $K = 1$ provided that Φ satisfies the RIP of order 2 with isometry constant $\delta < 1/3$.

Chapter 7

Sparse Recovery in White Noise

In practical settings such as those described in Chapter 5, there may be many sources that contribute to noise in our measurements, including noise present in the signal x, noise caused by the measurement hardware, quantization noise, and transmission errors in the case where the measurements are sent over a noisy channel. Fortunately, the RIP can provide us with a guarantee of stability to noise contaminating the measurements for many of the algorithms described in Section 2.4. In general, it can be shown that if $y = \Phi x + e$ with $x \in \Sigma_K$, then many common sparse recovery algorithms will yield a recovered signal $\widehat{x}$ satisfying

$$\|\widehat{x} - x\|_2 \leq C_0 \|e\|_2, \tag{7.1}$$

as described in Section 2.4. Thus, CS systems are stable in the sense that if the measurement error is bounded, then the reconstruction error is also bounded.

In this chapter[1] we analyze the impact of noise on the acquisition and recovery process more carefully. We first discuss the case where noise is added to the measurements, and examine the performance of an oracle-assisted recovery algorithm.

[1]This chapter builds on work done in collaboration with Richard G. Baraniuk and John Treichler [55]. Thanks also to J.P. Slavinsky for many useful discussions and helpful suggestions.

We conclude that the performance of most standard sparse recovery algorithms is near-optimal in the sense that it matches the performance of an oracle-assisted algorithm. We then consider the case where noise is added to the signal itself. In the case of white noise we show that compressive measurement systems like those described in Chapter 5 amplify this noise by an amount determined only by the number of measurements taken. Specifically, we observe that the recovered signal-to-noise ratio (SNR) decreases by 3dB each time the number of measurements is reduced by a factor of 2. This suggests that in low SNR settings, CS-based acquisition systems will be highly susceptible to noise.

7.1 Impact of Measurement Noise on an Oracle

To begin, we take a closer look at the problem of sparse signal recovery in the presence of measurement noise. Rather than directly analyzing a particular reconstruction algorithm, we will instead consider the performance of an oracle-assisted recovery algorithm that has perfect knowledge of the true location of the K nonzeros of x, which we denote $\Lambda = \text{supp}(x)$. While an oracle is typically not available, it characterizes the best that we can hope to achieve using any practical algorithm. In fact, we find that practical algorithms like CoSaMP typically perform almost as well as the oracle-assisted recovery algorithm.

Specifically, the oracle-assisted recovery algorithm is to solve

$$\widehat{x} = \arg\min_{x} \ \|\Phi x - y\|_2 \quad \text{subject to} \quad \text{supp}(x) = \Lambda, \tag{7.2}$$

where Λ is provided by an oracle. Recall from (6.1) that the least-squares optimal

recovery of x restricted to the index set Λ is given by

$$\begin{aligned} \widehat{x}|_\Lambda &= \Phi_\Lambda^\dagger y \\ \widehat{x}|_{\Lambda^c} &= 0. \end{aligned} \tag{7.3}$$

Before establishing our main result concerning oracle-assisted recovery, we first establish the following useful lemma. In the statement of the lemma, we use the notation $s_j(A)$ to denote the j^{th} nonzero singular value of A, i.e., $s_j(A)$ is the square root of the j^{th} eigenvalue of $A^T A$.

Lemma 7.1. *Suppose that Φ is an $M \times N$ matrix and let Λ be a set of indices with $|\Lambda| \leq K$ and $\{s_j(\Phi_\Lambda^\dagger)\}_{j=1}^K$ denote the K nonzero singular values of $\Phi_\Lambda^\dagger$. If Φ satisfies the RIP of order K with constant δ, then for $j = 1, 2, \ldots, K$ we have*

$$\frac{1}{\sqrt{1+\delta}} \leq s_j(\Phi_\Lambda^\dagger) \leq \frac{1}{\sqrt{1-\delta}}. \tag{7.4}$$

Proof. From the fact that Φ satisfies the RIP we immediately have that for any $u \in \mathbb{R}^K$,

$$(1-\delta)u^T u \leq u^T \Phi_\Lambda^T \Phi_\Lambda u \leq (1+\delta)u^T u,$$

and thus by picking u to be the K singular vectors, we have that

$$s_j(\Phi_\Lambda) \in \left[\sqrt{1-\delta}, \sqrt{1+\delta}\right]$$

for $j = 1, 2, \ldots, K$. Next recall that from the singular value decomposition (SVD) we can write

$$\Phi_\Lambda = U \Sigma V^T,$$

where U is an $M \times K$ matrix with orthonormal columns, V is a $K \times K$ unitary matrix, and Σ is a $K \times K$ diagonal matrix whose diagonal entries are the singular

values $s_j(\Phi_\Lambda)$.[2] Using this representation, and assuming that Φ_Λ is full rank, we can write

$$\begin{aligned}
\Phi_\Lambda^\dagger &= \left((U\Sigma V^T)^T U\Sigma V^T\right)^{-1} (U\Sigma V^T)^T \\
&= \left(V\Sigma U^T U\Sigma V^T\right)^{-1} V\Sigma U^T \\
&= \left(V\Sigma^2 V^T\right)^{-1} V\Sigma U^T \\
&= V\left(\Sigma^2\right)^{-1} V^T V\Sigma U^T \\
&= V\Sigma^{-1}U^T.
\end{aligned}$$

Thus, the SVD of $\Phi_\Lambda^\dagger$ is given by $V\Sigma^{-1}U^T$, and hence the singular values $s_j(\Phi_\Lambda^\dagger)$ are simply given by $1/s_j(\Phi_\Lambda)$, which establishes (7.4). □

This allows us to prove the following result.

Theorem 7.1. *Suppose that Φ satisfies the RIP of order K with constant δ. If $y = \Phi x + e$ where $x \in \Sigma_K$ and e is an arbitrary vector in $\mathbb{R}^M$, then the recovery provided by (7.3) when $\Lambda = \operatorname{supp}(x)$ satisfies*

$$\|\widehat{x} - x\|_2^2 \leq \frac{\|e\|_2^2}{1-\delta}. \tag{7.5}$$

Proof. We begin with the observation that when using the oracle, we have that $\widehat{x}|_{\Lambda^c} = x|_{\Lambda^c}$, so that

$$\begin{aligned}
\|\widehat{x} - x\|_2 = \|\widehat{x}|_\Lambda - x|_\Lambda\|_2 &= \|\Phi_\Lambda^\dagger y - x|_\Lambda\|_2 \\
&= \|(\Phi_\Lambda^T\Phi_\Lambda)^{-1}\Phi_\Lambda^T(\Phi x + e) - x|_\Lambda\|_2 \\
&= \|(\Phi_\Lambda^T\Phi_\Lambda)^{-1}\Phi_\Lambda^T(\Phi_\Lambda x|_\Lambda + e) - x|_\Lambda\|_2 \\
&= \|x|_\Lambda + (\Phi_\Lambda^T\Phi_\Lambda)^{-1}\Phi_\Lambda^T e - x|_\Lambda\|_2 = \|\Phi_\Lambda^\dagger e\|_2.
\end{aligned}$$

[2]Note that we are considering here the reduced form of the SVD.

In words, the oracle-assisted recovery algorithm will achieve an exact recovery of x, but the recovery will be contaminated by the noise $\|\Phi_\Lambda^\dagger e\|_2$. We can bound this error since, from Lemma 7.1, we have that the maximum singular value of $\Phi_\Lambda^\dagger$ is bounded above by $1/\sqrt{1-\delta}$. Thus, for any $e \in \mathbb{R}^M$, we have that

$$\|\Phi_\Lambda^\dagger e\|_2^2 \leq \frac{\|e\|_2^2}{1-\delta},$$

which establishes (7.5). □

Thus, the bound in (7.1) is optimal (up to a constant factor), since it matches the performance of an oracle-assisted recovery algorithm.

7.2 Impact of White Measurement Noise

While Theorem 7.1 characterizes the *worst-case* performance of the oracle-assisted recovery algorithm in the presence of *arbitrary* noise, it is also instructive to consider the *expected* performance in a more typical form of measurement noise. For example, in many common settings it is more natural to assume that the noise vector $e \sim \mathcal{N}(0, \sigma^2 I)$, i.e., e is i.i.d. Gaussian noise. We will consider the more general case where e is generated according to a *white* noise process, meaning that

$$\mathbb{E}(e) = 0 \tag{7.6}$$

and

$$\mathbb{E}\left(ee^T\right) = \sigma^2 I. \tag{7.7}$$

In other words, e is zero-mean, uncorrelated noise. Note that (7.7) implies that for any j, $\mathbb{E}(e_j^2) = \sigma^2$, and for any $j \neq i$, $\mathbb{E}(e_j e_i) = 0$. Thus, in this case we have that

$$\begin{aligned}\mathbb{E}\left(\|e\|_2^2\right) &= \mathbb{E}\left(\sum_{j=1}^{M} e_j^2\right) \\ &= \sum_{j=1}^{M} \mathbb{E}\left(e_j^2\right) = M\sigma^2.\end{aligned}$$

Hence, Theorem 7.1 might suggest that the best we can say is that given a typical noise vector, $\|\widehat{x} - x\|_2^2 \leq M\sigma^2/(1-\delta)$. However, we will now see that we actually can expect to do somewhat better than this.

Theorem 7.2. *Suppose that $y = \Phi x + e$ where $e \in \mathbb{R}^M$ is a white random vector satisfying (7.6) and (7.7). Furthermore, suppose that $x \in \Sigma_K$ and that Φ satisfies the RIP of order K with constant δ. Then the oracle-assisted recovery algorithm with solution defined by (7.3) for $\Lambda = \operatorname{supp}(x)$ satisfies*

$$\frac{K\sigma^2}{1+\delta} \leq \mathbb{E}\left(\|x - \widehat{x}\|_2^2\right) \leq \frac{K\sigma^2}{1-\delta}. \tag{7.8}$$

Proof. Recall that for the oracle-assisted recovery algorithm, we have that

$$\widehat{x}|_\Lambda = x|_\Lambda + \Phi_\Lambda^\dagger e.$$

Thus, our goal is to estimate $\mathbb{E}\left(\|\Phi_\Lambda^\dagger e\|_2^2\right)$. Towards this end, we first note that for any $K \times M$ matrix A with entries a_{ij}, since e is a white random vector, we have

$$\begin{aligned}\mathbb{E}\left(\|Ae\|_2^2\right) &= \mathbb{E}\left(\sum_{i=1}^{K} [Ae]_i^2\right) \\ &= \mathbb{E}\left(\sum_{i=1}^{K}\left(\sum_{j=1}^{M} a_{ij} e_j\right)^2\right)\end{aligned}$$

$$
\begin{aligned}
&= \mathbb{E}\left(\sum_{i=1}^{K}\left(\sum_{j=1}^{M} a_{ij}^2 e_j^2 + \sum_{j\neq k} a_{ij}e_j a_{ik}e_k\right)\right) \\
&= \sum_{i=1}^{K}\left(\sum_{j=1}^{M} a_{ij}^2 \mathbb{E}\left(e_j^2\right) + \sum_{j\neq k} a_{ij}a_{ik}\mathbb{E}\left(e_j e_k\right)\right) \\
&= \sum_{i=1}^{K}\sum_{j=1}^{M} a_{ij}^2\sigma^2 = \sigma^2\|A\|_F^2,
\end{aligned}
$$

where $\|\cdot\|_F$ denotes the Frobenius norm of A. Next we recall that the Frobenius norm of a $K \times M$ matrix with $K < M$ can also be calculated as

$$\|A\|_F^2 = \sum_{j=1}^{K} s_j(A)^2,$$

where $\{s_j(A)\}_{j=1}^{K}$ represent the singular values of A. Thus,

$$\mathbb{E}\left(\|\Phi_\Lambda^\dagger e\|_2^2\right) = \sigma^2\sum_{j=1}^{K} s_j(\Phi_\Lambda^\dagger)^2,$$

From Lemma 7.1 we have that $s_j(\Phi_\Lambda^\dagger) \in \left[1/\sqrt{1+\delta}, 1/\sqrt{1-\delta}\right]$ for $j = 1, 2, \ldots, K$, and hence

$$\frac{K}{1+\delta} \leq \sum_{j=1}^{K} s_j\left(\Phi_\Lambda^\dagger\right)^2 \leq \frac{K}{1-\delta},$$

which establishes (7.8). □

Note that while $\mathbb{E}\left(\|e\|_2^2\right) = M\sigma^2$, $\mathbb{E}\left(\|x - \widehat{x}\|_2^2\right) \approx K\sigma^2$. Thus, the expected energy in the error is lower than the predicted worst-case bound by a factor of K/M. This will prove significant in the following sections.

7.3 Impact of White Signal Noise

We now consider the case where the signal, as opposed to the measurements, are contaminated with noise. Thus, rather than the standard setting where $y = \Phi x + e$,

we now consider the case where

$$y = \Phi(x + n) = \Phi x + \Phi n. \tag{7.9}$$

This noise situation is subtly different from the standard setting because the noise added to the measurements has now has been acted upon by the matrix Φ, and so it is possible that Φn could be potentially rather large. Our chief interest here is to understand how Φ impacts the signal noise.

In order to simplify our analysis, we will make two assumptions concerning Φ: *(i)* the rows of Φ are orthogonal and *(ii)* each row of Φ has equal norm. While these assumptions are not necessary to ensure that Φ satisfies the RIP, both are rather intuitive. For example, it seems reasonable that if we wish to take as few measurements as possible, then each measurement should provide as much new information about the signal as possible, and thus requiring the rows of Φ to be orthogonal seems natural. Moreover, the second assumption can simply be interpreted as requiring that each measurement have "equal weight". Note that randomly generated Φ matrices will approximately satisfy these properties, and if Φ is an orthogonal projection, then it automatically satisfies these properties. Furthermore, these assumptions hold for both of the Φ matrices corresponding to the practical architectures described in Chapter 5.

These properties essentially ensure that if n is white noise, then Φn will be white noise as well, allowing us to more easily analyze the impact of white signal noise as quantified in the following theorem.

Theorem 7.3. *Suppose that Φ satisfies the RIP of order K with constant δ. Suppose furthermore that the rows of Φ are orthogonal and that each row of Φ has equal norm. If $n \in \mathbb{R}^N$ is a zero-mean, white random vector with $\mathbb{E}\left(nn^T\right) = \sigma^2 I$, then Φn is also*

a zero-mean, white random vector with $\mathbb{E}\left(\Phi n(\Phi n)^T\right) = \widetilde{\sigma}^2 I$*, where*

$$\frac{N}{M}\sigma^2(1-\delta) \leq \widetilde{\sigma}^2 \leq \frac{N}{M}\sigma^2(1+\delta). \tag{7.10}$$

Proof. We begin by noting that

$$\mathbb{E}\left([\Phi n]_i\right) = \mathbb{E}\left(\sum_{j=1}^{N} \phi_{ij} n_j\right) = \sum_{j=1}^{N} \phi_{ij}\mathbb{E}\left(n_j\right) = 0,$$

so that Φn is zero-mean, as desired. Hence, we now consider $\mathbb{E}\left(\left[\Phi n(\Phi n)^T\right]_{ij}\right)$. We begin by considering the diagonal entries for which $i = j$. In this case we have that

$$\begin{aligned}
\mathbb{E}\left(\left[\Phi n(\Phi n)^T\right]_{ii}\right) &= \mathbb{E}\left(\left(\sum_{k=1}^{N} \phi_{ik} n_k\right)^2\right) \\
&= \mathbb{E}\left(\sum_{k=1}^{N} \phi_{ik}^2 n_k^2 + \sum_{k\neq\ell} \phi_{ik}\phi_{i\ell} n_k n_\ell\right) \\
&= \sum_{k=1}^{N} \phi_{ik}^2 \mathbb{E}(n_k^2) + \sum_{k\neq\ell} \phi_{ik}\phi_{i\ell}\mathbb{E}(n_k n_\ell) \\
&= \sum_{k=1}^{N} \phi_{ik}^2 \sigma^2 = \|\phi_i\|_2^2 \sigma^2,
\end{aligned}$$

where ϕ_i represents the i^{th} row of Φ. Note that, by assumption, $\|\phi_i\|_2^2 = \|\phi_1\|_2^2$ for all i, so that

$$\mathbb{E}\left(\left[\Phi n(\Phi n)^T\right]_{ii}\right) = \|\phi_1\|_2^2 \sigma^2 \tag{7.11}$$

for all i.

Before we calculate $\|\phi_1\|_2^2$, we consider the off-diagonal case where $i \neq j$. In this case we have that

$$\mathbb{E}\left(\left[\Phi n(\Phi n)^T\right]_{ij}\right) = \mathbb{E}\left(\left(\sum_{k=1}^{N} \phi_{ik} n_k\right)\left(\sum_{\ell=1}^{N} \phi_{j\ell} n_\ell\right)\right)$$

$$
\begin{aligned}
&= \mathbb{E}\left(\sum_{k=1}^{N} \phi_{ik}\phi_{jk}n_k^2 + \sum_{k\neq\ell} \phi_{ik}\phi_{j\ell}n_k n_\ell\right) \\
&= \sum_{k=1}^{N} \phi_{ik}\phi_{jk}\mathbb{E}(n_k^2) + \sum_{k\neq\ell} \phi_{ik}\phi_{j\ell}\mathbb{E}(n_k n_\ell) \\
&= \sum_{k=1}^{N} \phi_{ik}\phi_{jk}\sigma^2 \\
&= 0,
\end{aligned}
$$

where the last equality follows from the assumption that the rows of Φ are orthogonal. Thus, $\mathbb{E}\left(\Phi n(\Phi n)^T\right)$ is the identity matrix scaled by $\widetilde{\sigma}^2 = \sigma^2\|\phi_1\|_2^2$.

It remains to show (7.10). We begin by applying the RIP to the set of 1-sparse binary vectors, from which we obtain that for any j,

$$(1-\delta) \leq \sum_{i=1}^{M} \phi_{ij}^2 \leq (1+\delta).$$

Thus,

$$(1-\delta)N \leq \|\Phi\|_F^2(1+\delta)N.$$

Since each row of Φ has equal norm, we must have that that $\|\phi_1\|_2^2 = \|\Phi\|_F^2/M$, and hence

$$(1-\delta)\frac{N}{M} \leq \|\phi_1\|_2^2 \leq (1+\delta)\frac{N}{M},$$

which when combined with (7.11) yields the desired result. □

Thus, while the oracle-assisted recovery procedure served to mildly attenuate white noise added to the measurements, when the noise is added to the signal itself it can be highly amplified by the measurement process when $M \ll N$. This is directly analogous to a classical phenomenon known as *noise folding*.

7.4 Noise Folding in CS

Theorem 7.3 tells us that the kinds of Φ matrices used in CS will amplify white noise by a factor of N/M. This makes sense intuitively, since we are projecting all of the noise in the N-dimensional input signal down into the M-dimensional measurements y, and all of the noise power must be preserved. In the literature, this effect is known as noise folding.

In order to quantify the impact of noise folding, we define the *input signal-to-noise ratio* (ISNR) and *output signal-to-noise ratio* (OSNR) as

$$\text{ISNR} = \frac{\|x\|_2^2}{\|(x+n)|_\Gamma - x\|_2^2} \tag{7.12}$$

and

$$\text{OSNR} = \frac{\|x\|_2^2}{\|\widehat{x} - x\|_2^2}, \tag{7.13}$$

where $\Gamma = \text{supp}(x)$ and $\widehat{x}$ is the output of the oracle-assisted recovery algorithm in (7.3) applied to $y = \Phi(x+n)$. The ISNR essentially measures the SNR for an oracle-assisted denoising algorithm that has access to the full signal $x+n$. Since the oracle knows which elements should be zero, it is able to achieve zero error on those coefficients — the only impact of the noise is on the nonzero coefficients. The OSNR measures the SNR for an oracle-assisted algorithm which must recover the original signal from the measurements $y = \Phi(x+n)$. We now define the *expected SNR loss* as

$$\text{Expected SNR loss} = \frac{\mathbb{E}(\text{ISNR})}{\mathbb{E}(\text{OSNR})} = \frac{\mathbb{E}\left(\|\widehat{x} - x\|_2^2\right)}{\mathbb{E}\left(\|(x+n)|_\Gamma - x\|_2^2\right)}. \tag{7.14}$$

In the event that the noise n is a white random vector, we can estimate the expected SNR loss as follows.

Theorem 7.4. *Suppose that Φ satisfies the RIP of order K with constant δ. Suppose furthermore that the rows of Φ are orthogonal and that each row of Φ has equal norm. If $n \in \mathbb{R}^N$ is a zero-mean, white random vector, then the expected SNR loss is bounded by*

$$\frac{N}{M} \cdot \frac{1-\delta}{1+\delta} \leq \text{Expected SNR loss} \leq \frac{N}{M} \cdot \frac{1+\delta}{1-\delta}. \tag{7.15}$$

Proof. Since n is white, we have that $\mathbb{E}\left(nn^T\right) = \sigma^2 I$. From this and the fact that $(x+n)|_\Gamma - x = n|_\Gamma$, we have that

$$\mathbb{E}\left(\|(x+n)|_\Gamma - x\|_2^2\right) = K\sigma^2. \tag{7.16}$$

We then observe that from Theorem 7.3, we have that $y = \Phi x + \Phi n$, where Φn is a white random vector with $\mathbb{E}\left(\Phi n(\Phi n)^T\right) = \widetilde{\sigma}^2 I$, where $\widetilde{\sigma}$ satisfies (7.10). Since Φn is white, we can apply Theorem 7.2 to obtain

$$\frac{K\widetilde{\sigma}^2}{1+\delta} \leq \mathbb{E}\left(\|x - \widehat{x}\|_2^2\right) \leq \frac{K\widetilde{\sigma}^2}{1-\delta}.$$

By combining this with the bound for $\widetilde{\sigma}$ in (7.10) we obtain

$$K\sigma^2 \cdot \frac{N}{M} \cdot \frac{1-\delta}{1+\delta} \leq \mathbb{E}\left(\|x - \widehat{x}\|_2^2\right) \leq K\sigma^2 \cdot \frac{N}{M} \cdot \frac{1+\delta}{1-\delta}. \tag{7.17}$$

Taking the ratio of (7.17) and (7.16) and simplifying establishes the theorem. □

Noise folding has a significant impact on the amount of noise present in CS measurements. Specifically, if we measure the expected SNR loss in dB, then we have that

$$\text{Expected SNR loss} \approx 10 \log_{10}\left(\frac{N}{M}\right).$$

Thus, every time we cut M in half (a one octave increase in the amount of subsampling), the expected SNR loss increases by 3dB. In other words, *for the acquisition of*

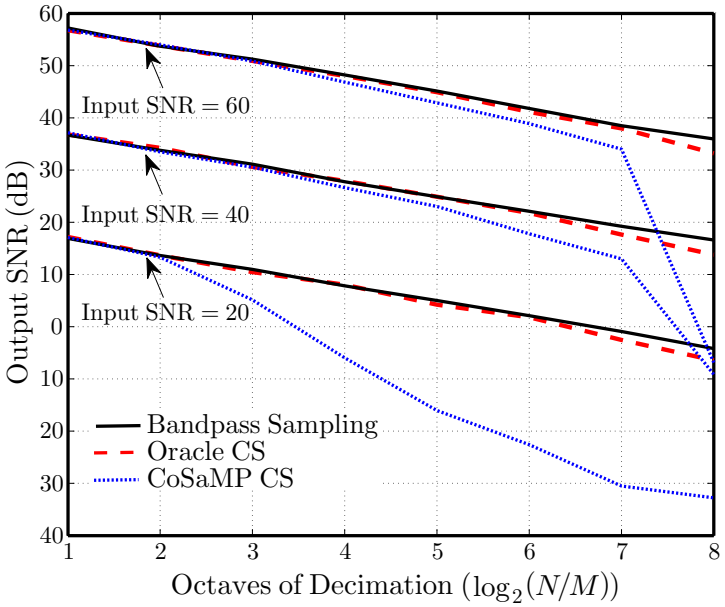

Figure 7.1: Simulation of signal recovery in noise. Output SNR as a function of the subsampling ratio N/M for a signal consisting of a single unmodulated voice channel in the presence of additive white noise.

a sparse signal in white noise, the SNR of the recovered signal decreases by 3 dB for every octave increase in the amount of subsampling.

We note that alternative signal acquisition techniques like *bandpass sampling* (sampling a narrowband signal uniformly at a sub-Nyquist rate to preserve the values but not the locations of its large Fourier coefficients) are affected by an identical 3dB/octave SNR degradation [120]. However, in practice bandpass sampling suffers from the limitation that it is impossible to determine the original original center frequencies after sampling. Furthermore, if there are multiple narrowband signals present, then bandpass sampling causes irreversible aliasing, in which case the components can overlap and will be impossible to separate. In contrast to bandpass sampling, however, CS acquisition preserves sufficient information to enable the recovery of both the values and the locations of the large Fourier coefficients.

The 3dB/octave SNR degradation represents an important tradeoff in the design of CS-based acquisition systems. Figure 7.1 shows the results of a set of simulations of a CS-based wideband signal acquisition system. In this case the signal to be

acquired consists of a single 3.1 kHz-wide unmodulated voice signal single-side-band-upconverted to a frequency within a 1 MHz input bandwidth of the receiver. In this case performance is measured as a function of the subsamping ratio N/M. The testing shown in Figure 7.1 was conducted at three input SNRs — 60, 40, and 20 dB — where input SNR in this case is simply the ratio of the signal power to that of the noise within the 3.1 kHz bandwidth occupied by the signal. The output SNR, measured classically within the 3.1 kHz signal bandwidth, was evaluated three ways:

- Bandpass sampling — This is not a recommended practical technique, but it serves as a benchmark since it is "filterless" like CS. It is important to note that this method "folds" the input spectrum so that signal frequencies can no longer be unambiguously determined at the receiver.

- Oracle-assisted signal recovery from compressive measurements — While not practical, again, the oracle provides a way to determine what portion of any observed received quality degradation is totally unavoidable within the CS framework and what portion is due to the recovery algorithm's inability to determine the spectral support.

- Practical CS-based signal recovery using CoSaMP to determine the spectral support of the input signal.

We can make several observations from the experimental results depicted in Figure 7.1. First, we note that for small amounts of subsampling the output SNR of both the bandpass sampled signal and the oracle-assisted CS recovery is degraded at a rate of 3dB for each octave increase in the ratio N/M, exactly as predicted by theory. Next, we note that the output SNR of the oracle-assisted recovery approach closely follows the bandpass sampling output SNR across the entire range considered for N/M. The performance of the CoSaMP algorithm generally tracks the others, but performs progressively more poorly for high subsampling ratios. Moreover, its

performance collapses as the theoretical limit is reached and as the input SNR falls below a critical level. Specifically, our theory requires that $M > C_1 K \log(N/K)$, and thus we must have that

$$\frac{N}{M} < \frac{1}{C_1} \frac{(N/K)}{\log(N/K)}.$$

Note that for these experiments, $N/K = (2 \cdot 10^6)/(3.1 \cdot 10^3) \approx 645$, and thus, ignoring the effect of the unknown constant C_1, we should expect that the maximum allowable amount of subsampling should be bounded roughly by $645/\log(645) \approx 100$. This corresponds to $\log_2(N/M) \approx 6.6$. In Figure 7.1 we observe that we do not begin to observe a dramatic difference between the performance of oracle-assisted CS and CoSaMP until $\log_2(N/M) > 7$. In the regimes where the performance of CoSaMP is significantly worse than that of oracle-assisted recovery, we observe that oracle-assisted recovery continues to match the SNR of the bandpass sampled signal. This indicates that in these regimes, CoSaMP is unable to identify the correct locations of the nonzero Fourier coefficients, since if it could it would match the oracle-assisted recovery approach, i.e., support estimation is the harder part of CS recovery (as opposed to coefficient estimation). Thus, if any side information concerning the likely locations of these nonzeros were available, then one could expect that exploiting this information would have a significant impact on the SNR performance.

Chapter 8

Sparse Recovery in Sparse Noise

In Chapter 7 we considered the case where our signal or measurements were corrupted with unstructured noise that was either bounded or bounded with high probability. These results are well-suited to deal with noise that is evenly distributed across the signal or measurements, such as i.i.d. Gaussian, thermal, or quantization noise. However, in other cases our noise will satisfy some additional structure. We will have more to say regarding structured signal noise in Chapter 10, but in this chapter[1] we analyze the case where the noise itself is also sparse. We demonstrate that in addition to satisfying the RIP, the same random matrices considered in Chapter 4 satisfy an additional property that leads to measurements that are guaranteed to be robust to a small number of arbitrary corruptions and to other forms of sparse measurement noise. We propose an algorithm dubbed *Justice Pursuit* that can exploit this structure to recover sparse signals in the presence of corruption. We then show that this structure can be viewed as an example of a more general phenomenon. Specifically, we propose a definition of *democracy* in the context of CS and leverage our analysis of Justice Pursuit to show that random measurements are democratic. We conclude with a brief discussion of the broader role of democracy in CS.

[1] This work was done in collaboration with Richard G. Baraniuk, Petros T. Boufounos, and Jason N. Laska [121, 122].

8.1 Measurement Corruption

In this chapter, we consider a more structured measurement noise model, namely

$$y = \Phi x + \Omega e, \tag{8.1}$$

where Ω is an $M \times L$ matrix with $L \leq M$ orthonormal columns, and the vector e is sparse. The matrix Ω represents the basis or subspace in which the noise is sparse. The case where $\Omega = I$ is representative of many practical sources of noise. For example, there may be short bursts of high noise, or certain measurements may be invalid because of defective hardware or spikes in the power supply. When measurements are sent over a network, some measurements may be lost altogether, or in a sensor network, malfunctioning sensors may regularly transmit corrupted measurements while the other sensors do not. In these cases the noise is sparse in the canonical basis. In other settings, the measurement noise may be sparse or compressible when represented in some transform basis. For example, the measurements could be corrupted with 60Hz hum,[2] in which case the noise is sparse in the Fourier basis. Similarly, measurement noise from a DC bias that changes abruptly would be piecewise-smooth and thus sparse in a wavelet basis.

In these cases, $\|e\|_2$ may be extremely large, and thus the resulting bound $C_0\|e\|_2$ on the reconstruction error will also be large. However, one can hope to do much better. To see why, suppose that the measurement noise is sparse in the basis I so that only a few of the measurements are corrupted with large errors and that the remaining measurements are noise-free. Standard recovery algorithms will return a signal estimate $\widehat{x}$ that satisfies only $\|\widehat{x} - x\|_2 \leq C_0\|e\|_2$. However, if we knew which measurements were corrupted, then we could simply ignore them. If Φ is generated randomly with M sufficiently large, and if the locations of the corrupted

[2] In some regions hum consists of a 50Hz sinusoid (and its harmonics).

measurements were known *a priori*, then the signal could be reconstructed exactly by using only the noiseless measurements [123]. The challenge is that it is typically not possible to know exactly which measurements have been corrupted.

8.2 Justice Pursuit

Our goal is to design an algorithm that will recover both the signal and noise vectors by leveraging their sparsity. Towards this end, suppose that we acquire measurements of the form in (8.1) and that $x \in \Sigma_K$ and $e \in \Sigma_\kappa$. Note that the measurements can be expressed in terms of an $M \times (N + L)$ matrix multiplied by a $(K + \kappa)$-sparse vector:

$$\Phi x + \Omega e = [\Phi \ \Omega] \begin{bmatrix} x \\ e \end{bmatrix}. \tag{8.2}$$

We now introduce our reconstruction program, *Justice Pursuit* (JP):

$$\widehat{u} = \arg\min_u \|u\|_1 \quad \text{subject to} \quad [\Phi \ \Omega]u = y, \tag{8.3}$$

where $\widehat{u}$ is an intermediate $(N + L) \times 1$ recovery vector. The signal estimate $\widehat{x}$ is obtained by selecting the first N elements of $\widehat{u}$, i.e., $\widehat{x}_i = \widehat{u}_i, \ i = 1, \ldots, N$. Furthermore, an estimate of the noise vector $\widehat{e}$ can be obtained by selecting the last L elements of $\widehat{u}$, i.e., $\widehat{e}_i = \widehat{u}_{i+N}, \ i = 1, \ldots, L$. Note that one can also adapt any of the iterative algorithms from Sections 2.4.2 and 2.4.3 by simply replacing Φ with $[\Phi \ \Omega]$.

JP is essentially identical to a program proposed independently in [124, 125]. Note, however, that in [124, 125] the authors consider only Φ that are composed of a set of highly correlated training vectors and do not consider this program within the more traditional context of CS. Indeed, due to our differing assumptions on Φ, we can demonstrate stronger, non-asymptotic guarantees on the recovery of x and e provided by JP. The sparse noise model has also been considered in the context of

CS in [126]; however the authors use a probabilistic approach for the analysis, a specialized measurement scheme, and propose a non-convex program with non-linear constraints for signal recovery, resulting in substantial differences from the results we present below. Note also that while [77] also considers the use of ℓ_1-minimization to mitigate sparse noise, this is in the context of error correction coding. In this framework the signal to be encoded is not necessarily sparse and $M > N$, resulting in a substantially different approach.

While JP is relatively intuitive, it is not clear that it will necessarily work. In particular, in order to analyze JP using standard methods, we must show that the matrix $[\Phi\ \Omega]$ satisfies the RIP. We now demonstrate that for any choice of Ω, if we draw the entries of Φ according to a sub-Gaussian distribution as in Chapter 4, then with high probability $[\Phi\ \Omega]$ will satisfy the RIP for any Ω. To do so, we use several results from Chapter 4 to establish the following lemma, which demonstrates that for any u, if we draw Φ at random, then $\|[\Phi\ \Omega]u\|_2$ is concentrated around $\|u\|_2$.

Lemma 8.1. *Suppose that Φ is an $M \times M$ matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$ and let Ω be an $M \times L$ matrix with orthonormal columns. Furthermore, let $u \in \mathbb{R}^{N+L}$ be an arbitrary vector with the first N entries denoted by x and the last L entries denoted by e. Then for any $\epsilon > 0$, and any $u \in \mathbb{R}^{N+L}$,*

$$\mathbb{E}\left(\|[\Phi\ \Omega]u\|_2^2\right) = \|u\|_2^2 \tag{8.4}$$

and

$$\mathbb{P}\left(\left|\|[\Phi\ \Omega]u\|_2^2 - \|u\|_2^2\right| \geq \epsilon\|u\|_2^2\right) \leq 4e^{-M\epsilon^2/32}. \tag{8.5}$$

Proof. We first note that since $[\Phi\ \Omega]u = \Phi x + \Omega e$,

$$\begin{aligned}\|[\Phi\ \Omega]u\|_2^2 &= \|\Phi x + \Omega e\|_2^2 \\ &= (\Phi x + \Omega e)^T(\Phi x + \Omega e)\end{aligned}$$

$$= x^T\Phi^T\Phi x + 2e^T\Omega^T\Phi x + e^T\Omega^T\Omega e$$
$$= \|\Phi x\|_2^2 + 2e^T\Omega^T\Phi x + \|e\|_2^2. \tag{8.6}$$

From Corollary 4.2 we have that $\mathbb{E}\left(\|\Phi x\|_2^2\right) = \|x\|_2^2$. Furthermore, using Lemma 4.2 it is straightforward to show that $2e^T\Omega^T\Phi x \sim \text{SSub}\left(4\|x\|_2^2\|\Omega e\|_2^2/M\right)$, since the elements of Φx are strictly sub-Gaussian variables with variance $\|x\|_2^2/M$. Thus, from Lemma 4.1 we have that $\mathbb{E}\left(2e^T\Omega^T\Phi x\right) = 0$. Hence, from (8.6) we obtain

$$\mathbb{E}\left(\|[\Phi\ \Omega]u\|_2^2\right) = \|x\|_2^2 + \|e\|_2^2,$$

and since $\|u\|_2^2 = \|x\|_2^2 + \|e\|_2^2$, this establishes (8.4).

We now turn to (8.5). From Corollary 4.2

$$\mathbb{P}\left(\left|\|Y\|_2^2 - \|x\|_2^2\right| \geq \delta\|x\|_2^2\right) \leq 2\exp\left(-\frac{M\delta^2}{C^*}\right). \tag{8.7}$$

As noted above, $2e^T\Omega^T\Phi x \sim \text{SSub}\left(4\|x\|_2^2\|\Omega e\|_2^2/M\right)$. Note that since the columns of Ω are orthonormal, $\|\Omega e\|_2^2 = \|e\|_2^2$. Hence, from Theorem 4.1 we have that

$$\mathbb{P}\left(\left|2e^T\Omega^T\Phi x\right| \geq \delta\|x\|_2\|e\|_2\right) \leq 2e^{-M\delta^2/8}. \tag{8.8}$$

Thus, since $C^* \approx 6.52 < 8$, we can combine (8.7) and (8.8) to obtain that with probability at least $1 - 4e^{-M\delta^2/8}$ we have that both

$$(1-\delta)\|x\|_2^2 \leq \|\Phi x\|_2^2 \leq (1+\delta)\|x\|_2^2 \tag{8.9}$$

and

$$-\delta\|x\|_2\|e\|_2 \leq 2e^T\Omega^T\Phi x \leq \delta\|x\|_2\|e\|_2. \tag{8.10}$$

Using (8.6), we can combine (8.9) and (8.10) to obtain

$$\begin{aligned}\|[\Phi\ \Omega]u\|_2^2 &\leq (1+\delta)\|x\|_2^2 + \delta\|x\|_2\|e\|_2 + \|e\|_2^2 \\ &\leq (1+\delta)\left(\|x\|_2^2 + \|e\|_2^2\right) + \delta\|x\|_2\|e\|_2 \\ &\leq (1+\delta)\|u\|_2^2 + \delta\|u\|_2^2 \\ &= (1+2\delta)\|u\|_2^2,\end{aligned}$$

where the last inequality follows from the fact that $\|x\|_2\|e\|_2 \leq \|u\|_2\|u\|_2$. Similarly, we also have that

$$\|[\Phi\ \Omega]u\|_2^2 \geq (1-2\delta)\|u\|_2^2.$$

By substituting $\epsilon = \delta/2$, this establishes (8.5). □

Using Lemma 8.1, we now demonstrate that the matrix $[\Phi\ \Omega]$ satisfies the RIP provided that M is sufficiently large. This theorem follows immediately from Lemma 8.1 by using a proof identical to that of Theorem 4.3, so we omit the proof for the sake of brevity.

Theorem 8.1. *Fix $\delta \in (0,1)$. Let Φ be an $M \times N$ random matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$ and let Ω be an $M \times L$ matrix with orthonormal columns. If*

$$M \geq C_1(K+\kappa)\log\left(\frac{N+L}{K+\kappa}\right), \tag{8.11}$$

then $[\Phi\ \Omega]$ satisfies the RIP of order $(K+\kappa)$ with the prescribed δ with probability exceeding $1 - 4e^{-C_2 M}$, where C_1 is arbitrary and $C_2 = \delta^2/64 - \log(42e/\delta)/C_1$.

Theorem 8.1 implies that when both x and e are sparse, JP recovers both x and e exactly. Thus, even if $\|e\|_2$ is unbounded, in this setting JP achieves optimal performance. To summarize, the hallmarks of JP include:

1. *exact recovery* of the sparse signal x;

2. *exact recovery* of the sparse noise term e;

3. *blindness* to the locations and size of the measurement errors — thus, the corrupted measurements could be adversarially selected and the noise on the corrupted measurements can be arbitrarily large;

4. no user-defined parameters;

5. standard CS recovery algorithm implementations can be trivially modified, i.e., *justified*, to perform JP, so that optimized routines can be easily adapted to this setting.

In the case where e contains additional sources of noise that are not sparse, e.g., AWGN or quantization error in addition to hum, but has norm bounded by ϵ, we propose an algorithm we dub *Justice Pursuit De-Noising* (JPDN):

$$\widehat{u} = \arg\min_{u} \|u\|_1 \quad \text{s.t.} \quad \|[\Phi\ \Omega]u - y\|_2 < \epsilon. \tag{8.12}$$

The performance guarantees of JPDN are analogous to those for BPDN. Specifically, from Theorem 3.2 we have that provided $[\Phi\ \Omega]$ satisfies the RIP of order $K+\kappa$ with constant δ sufficiently small, we have

$$\|\widehat{u} - u\|_2 \leq C_1 \frac{\sigma_{K+\kappa}(u)_1}{\sqrt{K+\kappa}}. \tag{8.13}$$

Note that we trivially have that $\|\widehat{x} - x\|_2 \leq \|\widehat{u} - u\|_2$, and since one possible $K+\kappa$-sparse approximation to u consists of taking the K largest coefficients of x and the κ largest coefficients of e, we also have that $\sigma_{K+\kappa}(u)_1 \leq \sigma_K(x)_1 + \sigma_\kappa(e)_1$. Thus, from (8.13) we also have

$$\|\widehat{x} - x\|_2 \leq C_1 \frac{\sigma_K(x)_1 + \sigma_\kappa(e)_1}{\sqrt{K+\kappa}}.$$

This guarantees a degree of robustness to non-sparse noise or signals.

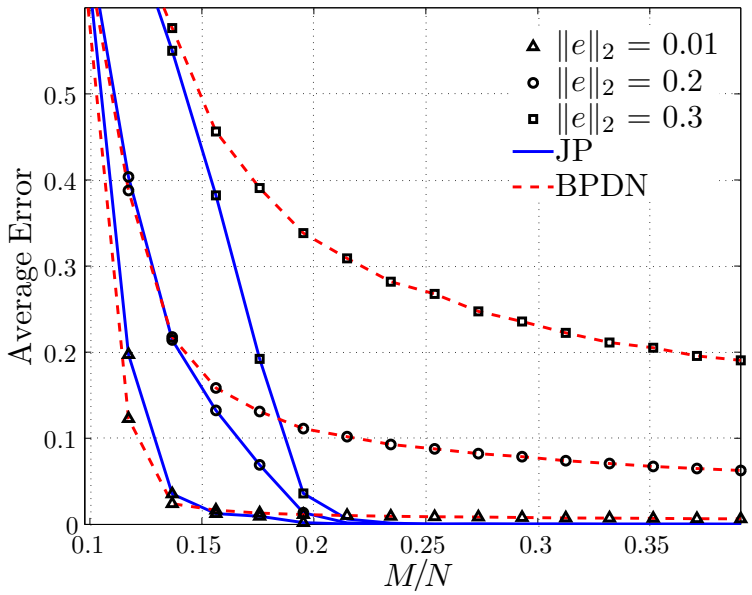

Figure 8.1: Comparison of average reconstruction error $\|x - \widehat{x}\|_2$ between JP and BPDN for noise norms $\|e\|_2 = 0.01,\ 0.2$, and 0.3. All trials used parameters $N = 2048$, $K = 10$, and $\kappa = 10$. This plot demonstrates that while BPDN never achieves exact reconstruction, JP does.

8.3 Simulations

8.3.1 Average performance comparison

In Figures 8.1 and 8.2, we compare the average reconstruction error of JP (solid lines) against the average error of BPDN (dashed lines). We perform two experiments, each with parameters $N = 2048$, $K = 10$, and $\|x\|_2 = 1$, with $M/N \in [0.1, 0.4]$, and record the average error $\|x - \widehat{x}\|_2$ over 100 trials.

In the first experiment, depicted in Figure 8.1, we fix $\|e\|_0 = \kappa = 10$ and vary $\|e\|_2$. We observe that the reconstruction error for BPDN does not decay to zero no matter how large we set M. Most representative of this is the $\|e\|_2 = 0.01$ case. As M/N increases, this line reaches a minimum value greater than zero and does not decay further. In contrast, JP reaches exact recovery in all tests. In the second experiment, depicted in Figure 8.2, we fix $\|e\|_2 = 0.1$ and vary κ. Again, the performance of BPDN does not decay to zero, and furthermore, the performance does not vary with κ on

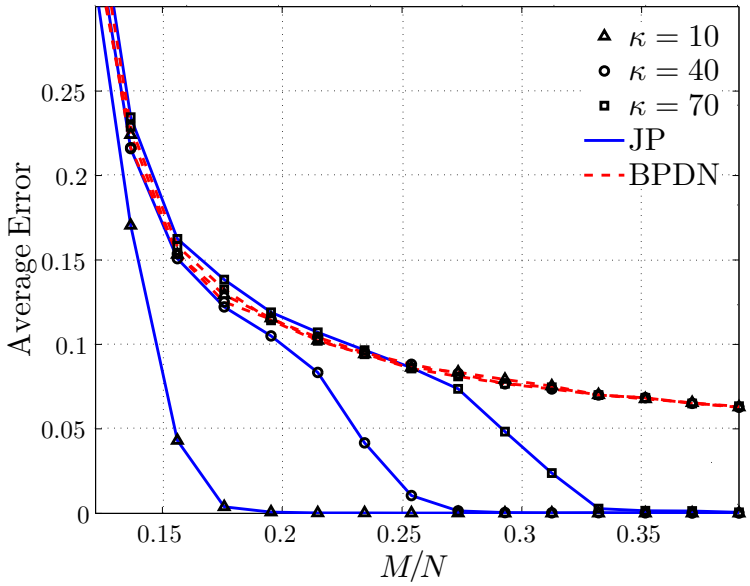

Figure 8.2: Comparison of average reconstruction error $\|x - \widehat{x}\|_2$ between JP and BPDN for $\kappa = 10,\ 40$, and 70. All trials used parameters $N = 2048$, $K = 10$, and $\|e\|_2 = 0.1$. This plot demonstrates that JP performs similarly to BPDN until M is large enough to reconstruct κ noise entries.

average. As expected the error of JP goes to zero and requires more measurements to do so as κ increases.

8.3.2 Reconstruction with hum

In this experiment we study the reconstruction performance from measurements corrupted by hum, meaning that we add a 60Hz sinusoid to the measurements. We use a 256×256 pixel test image that is compressible in the wavelet domain, set the measurement ratio to $M/N = 0.2$, and set the measurement signal-to-noise ratio (SNR) to 9.3dB, where measurement SNR in dB is defined as $10 \log_{10}(\|\Phi x\|_2^2 / \|e\|_2^2)$. We recover using BPDN with $\epsilon = \|e\|_2$ and using JP with the Fourier basis for Ω. Note that rather than choosing the entire Fourier basis, a matrix containing the 60Hz tone and its harmonics can be chosen to reduce the number of required measurements.

Figure 8.3(a) depicts the reconstruction from BPDN and Figure 8.3(b) depicts the reconstruction from JP. Both images contain compression artifacts, such as "ringing,"

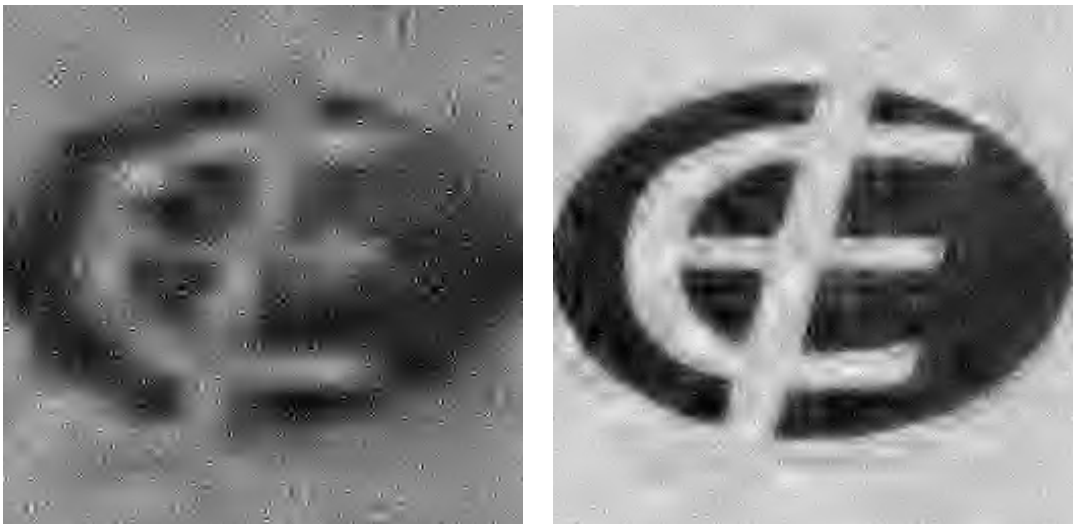

(a) (b)

Figure 8.3: Reconstruction of an image from CS measurements that have been distorted by an additive 60Hz sinusoid (hum). The experimental parameters are $M/N = 0.2$ and measurement SNR = 9.3dB. (a) Reconstruction using BPDN. (b) Reconstruction using JP. Spurious artifacts due to noise are present in the image in (a) but not in (b). Significant edge detail is lost in (a) but recovered in (b).

since the signal is not strictly sparse. However, the BPDN reconstruction contains spurious artifacts, due not to compression but to noise, while the JP reconstruction does not. Furthermore, significant edge detail is lost in the BPDN reconstruction.

8.3.3 Measurement denoising

In this experiment we use our algorithm to denoise measurements y that have been acquired by the single-pixel camera [56]. The image dimensions are 256×256 pixels and $M/N = 0.1$. The denoising procedure is as follows. First we reconstruct the image using JP with the Fourier basis for Ω. Second, because the measurement noise is not strictly sparse, we select the 15 largest terms from $\widehat{e}$, denoted as $\widehat{e}'$, and subtract their contribution from the original measurements, i.e.,

$$y' = y - \Omega\widehat{e}'.$$

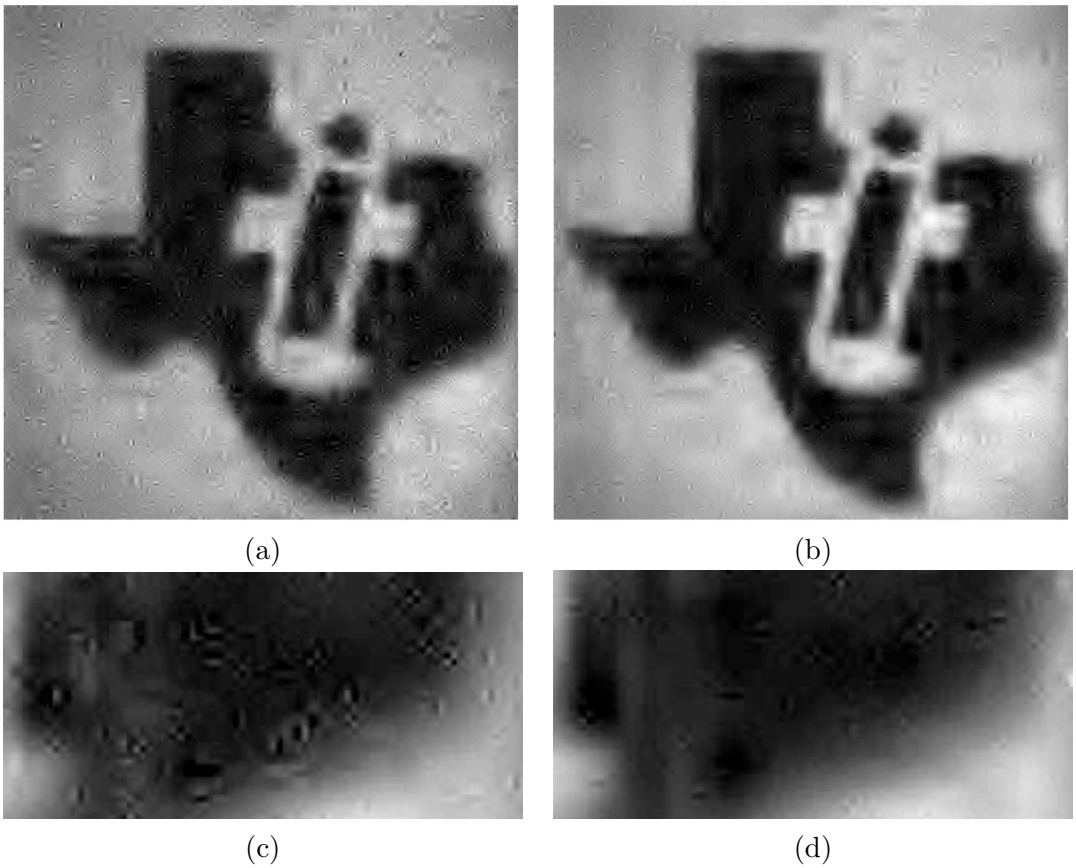

Figure 8.4: Reconstruction from CS camera data. (a) Reconstruction from CS camera measurements. (b) Reconstruction from denoised CS camera measurements. (c) and (d) depict zoomed sections of (a) and (b), respectively. Noise artifacts are removed without further smoothing of the underlying image.

Third, reconstruction from y' is performed with BPDN using the parameter $\epsilon = 0.3$. To compare, we also reconstruct the image from the original measurements y using BPDN with the same ϵ. In general, this procedure can be performed iteratively, selecting several spikes from $\widehat{e}$ at each iteration and subtracting their contribution from the measurements.

Figure 8.4(a) depicts the reconstruction from y and Figure 8.4(b) depicts the reconstruction from y', and Figures 8.4(c) and 8.4(d) show a zoomed section of each, respectively. The reconstruction from the original measurements contains significantly more spurious artifacts, while the reconstruction from denoised measurements removes these artifacts without further smoothing of the underlying image.

There are many topics that have not been fully explored in this section. For instance, the noise could be compressible rather than strictly sparse, or could consist of low energy noise on all measurements in addition to the sparse noise. For example, measurements may be subject to both shot noise and quantization errors simultaneously. Additionally, models can be employed to exploit additional noise structure and reduce the number of required measurements, or recover the signal with higher accuracy. Finally, the performance of JPDN or adaptations of greedy or iterative methods to this setting remains a topic of ongoing work.

8.4 Justice and Democracy

The moral of the preceding sections is that random measurements are *just*, meaning that they are robust to a small number of arbitrary corruptions. In this section, we investigate a closely related property of random measurements. Specifically, we show that random matrices are *democratic*, which has historically been taken to mean that each measurement carries roughly the same amount of information about the signal. We adopt a more precise definition, and further demonstrate that random measurements are robust to the *loss* of a small number of arbitrary measurements by building on the Lemma 8.1. In addition, we draw connections to oversampling and demonstrate stability from the loss of significantly more measurements.

8.4.1 Democracy

While it is not usually rigorously defined in the literature, democracy is usually taken to mean that each measurement contributes a similar amount of information about the signal x to the compressed representation y [127–129].[3] Others have de-

[3]The original introduction of this term was with respect to quantization [127, 128], i.e., a democratic quantizer would ensure that each bit is given "equal weight." As the CS framework developed, it became empirically clear that CS systems exhibited this property with respect to compression [129].

scribed democracy to mean that each measurement is equally important (or unimportant) [130]. Despite the fact that democracy is so frequently touted as an advantage of random measurements, it has received little analytical attention in the CS context. Perhaps more surprisingly, the property has not been explicitly exploited in applications until recently [123].

The fact that random measurements are democratic seems intuitive; when using random measurements, each measurement is a randomly weighted sum of a large fraction (or all) of the entries of x, and since the weights are chosen independently at random, no preference is given to any particular entries. More concretely, suppose that the measurements $y_1, y_2, \ldots, y_M$ are i.i.d. according to some distribution f_Y, as is the case for Φ with i.i.d. entries. Now suppose that we select $\widetilde{M} < M$ of the y_i at random (or according to some procedure that is *independent* of y). Then clearly, we are left with a length-$\widetilde{M}$ measurement vector $\widetilde{y}$ such that each $\widetilde{y}_i \sim f_Y$. Stated another way, if we set $D = M - \widetilde{M}$, then there is no difference between collecting $\widetilde{M}$ measurements and collecting M measurements and deleting D of them, provided that this deletion is done independently of the actual values of y.

However, following this line of reasoning will ultimately lead to a rather weak definition of democracy. To see this, consider the case where the measurements are deleted by an adversary. By adaptively deleting the entries of y one can change the distribution of $\widetilde{y}$. For example, the adversary can delete the D largest elements of y, thereby skewing the distribution of $\widetilde{y}$. In many cases, especially if the same matrix Φ will be used repeatedly with different measurements being deleted each time, it would be far better to know that *any* $\widetilde{M}$ measurements will be sufficient to reconstruct the signal. This is a significantly stronger requirement.

8.4.2 Democracy and the RIP

The RIP also provides us with a way to quantify our notion of democracy in the deterministic setting of CS. To do so, we first formally define democracy. In our definition, we assume that Φ is an $M \times N$ matrix and in the case where $\Gamma \subset \{1, 2, , \ldots, M\}$ we use the notation Φ^Γ to denote the $|\Gamma| \times N$ matrix obtained by selecting the rows of Φ indexed by Γ.

Definition 8.1. Let Φ be and $M \times N$ matrix, and let $\widetilde{M} \leq M$ be given. We say that Φ is $(\widetilde{M}, K, \delta)$-*democratic* if for all Γ such that $|\Gamma| \geq \widetilde{M}$ the matrix Φ^Γ satisfies the RIP of order K with constant δ.

We now show that sub-Gaussian matrices satisfy this property with high probability.

Theorem 8.2. *Fix $\delta \in (0, 1)$. Let Φ be an $M \times N$ random matrix whose entries ϕ_{ij} are i.i.d. with $\phi_{ij} \sim \mathrm{SSub}(1/M)$. Let $\widetilde{M} \leq M$ be given, and define $D = M - \widetilde{M}$. If*

$$M = C_1(K + D) \log\left(\frac{N + M}{K + D}\right), \tag{8.14}$$

then with probability exceeding $1 - 4e^{-C_2 M}$ we have that Φ is $(\widetilde{M}, K, \delta/(1 - \delta))$-democratic, where C_1 is arbitrary and $C_2 = \delta^2/64 - \log(42e/\delta)/C_1$.

Proof. Our proof consists of two main steps. We begin by defining the $M \times (N + M)$ matrix $\widetilde{\Phi} = [I\ \Phi]$ formed by appending Φ to the $M \times M$ identity matrix. Theorem 8.1 demonstrates that under the assumptions in the theorem statement, with probability exceeding $1 - 4e^{-C_2 M}$ we have that $\widetilde{\Phi}$ satisfies the RIP of order $K + D$ with constant δ. The second step is to use this fact to show that all possible $\widetilde{M} \times N$ submatrices of Φ satisfy the RIP of order K with constant $\delta/(1 - \delta)$.

Towards this end, we let $\Gamma \subset \{1, 2, \ldots, M\}$ be an arbitrary subset of rows such

that $|\Gamma| \geq \widetilde{M}$. Define $\Lambda = \{1, 2, \ldots, M\} \setminus \Gamma$ and note that $|\Lambda| = D$. Recall that

$$P_\Lambda = \widetilde{\Phi}_\Lambda \widetilde{\Phi}_\Lambda^\dagger, \tag{8.15}$$

defines the orthogonal projection onto $\mathcal{R}(\widetilde{\Phi}_\Lambda)$, i.e., the range, or column space, of $\widetilde{\Phi}_\Lambda$. Furthermore, we define

$$P_\Lambda^\perp \triangleq I - P_\Lambda, \tag{8.16}$$

as the orthogonal projection onto the orthogonal complement of $\mathcal{R}(\widetilde{\Phi}_\Lambda)$. In words, this projector annihilates the columns of $\widetilde{\Phi}$ corresponding to the index set Λ. Now, note that $\Lambda \subset \{1, 2, \ldots, M\}$, so $\widetilde{\Phi}_\Lambda = I_\Lambda$. Thus,

$$P_\Lambda = I_\Lambda I_\Lambda^\dagger = I_\Lambda (I_\Lambda^T I_\Lambda)^{-1} I_\Lambda^T = I_\Lambda I_\Lambda^T = I(\Lambda),$$

where we use $I(\Lambda)$ to denote the $M \times M$ matrix with all zeros except for ones on the diagonal entries corresponding to the columns indexed by Λ. (We distinguish the $M \times M$ matrix $I(\Lambda)$ from the $M \times D$ matrix I_Λ — in the former case we replace columns not indexed by Λ with zero columns, while in the latter we remove these columns to form a smaller matrix.) Similarly, we have

$$P_\Lambda^\perp = I - P_\Lambda = I(\Gamma).$$

Thus, we observe that the matrix $P_\Lambda^\perp \widetilde{\Phi} = I(\Gamma)\widetilde{\Phi}$ is simply the matrix $\widetilde{\Phi}$ with zeros replacing all entries on any row i such that $i \notin \Gamma$, i.e., $(P_\Lambda^\perp \widetilde{\Phi})^\Gamma = \widetilde{\Phi}^\Gamma$ and $(P_\Lambda^\perp \widetilde{\Phi})^\Lambda = 0$. Furthermore, Lemma 6.2 states that for $\widetilde{\Phi}$ satisfying the RIP of order $K + D$ with constant δ, we have that

$$\left(1 - \frac{\delta}{1-\delta}\right) \|u\|_2^2 \leq \|P_\Lambda^\perp \widetilde{\Phi} u\|_2^2 \leq (1+\delta)\|u\|_2^2, \tag{8.17}$$

holds for all $u \in \mathbb{R}^{N+M}$ such that $\|u\|_0 = K + D - |\Lambda| = K$ and $\text{supp}(u) \cap \Lambda = \emptyset$. Equivalently, letting $\Lambda^c = \{1, 2, \ldots, N+M\} \setminus \Lambda$, this result states that $(I(\Gamma)\widetilde{\Phi})_{\Lambda^c}$ satisfies the RIP of order K with constant $\delta/(1-\delta)$. To complete the proof, we note that if $(I(\Gamma)\widetilde{\Phi})_{\Lambda^c}$ satisfies the RIP of order K with constant $\delta/(1-\delta)$, then we trivially have that $I(\Gamma)\Phi$ also has the RIP of order at least K with constant $\delta/(1-\delta)$, since $I(\Gamma)\Phi$ is just a submatrix of $(I(\Gamma)\widetilde{\Phi})_{\Lambda^c}$. Since $\|I(\Gamma)\Phi x\|_2 = \|\Phi^\Gamma x\|_2$, this establishes the theorem. □

8.4.3 Robustness and stability

Observe that we require $O(D\log(N))$ additional measurements to ensure that Φ is $(\widetilde{M}, K, \delta)$-democratic compared to the number of measurements required to simply ensure that Φ satisfies the RIP of order K. This seems intuitive; if we wish to be robust to the loss of any D measurements while retaining the RIP of order K, then we should expect to take *at least* D additional measurements. This is not unique to the CS framework. For instance, by *oversampling*, i.e., sampling faster than the minimum required Nyquist rate, uniform sampling systems can also improve robustness with respect to the loss of measurements. Reconstruction can be performed in principle on the remaining non-uniform grid, as long as the remaining samples satisfy the Nyquist range on average [131].

However, linear reconstruction in such cases is known to be unstable. Furthermore the linear reconstruction kernels are difficult to compute. Under certain conditions stable non-linear reconstruction is possible, although this poses further requirements on the subset of samples that can be lost and the computation can be expensive [132]. For example, deleting contiguous groups of measurements can be a challenge for the stability of the reconstruction algorithms. Instead, the democratic property of random measurements allows for the deletion of an arbitrary subset D of the measurements without compromising the reconstruction stability, independent of the way

these measurements are chosen.

In some applications, this difference may have significant impact. For example, in finite dynamic range quantizers, the measurements saturate when their magnitude exceeds some level. Thus, when uniformly sampling with a low saturation level, if one sample saturates, then the likelihood that any of the neighboring samples will saturate is high, and significant oversampling may be required to ensure any benefit. However, in CS, if many adjacent measurements were to saturate, then for only a slight increase in the number of measurements we can mitigate this kind of error by simply rejecting the saturated measurements; the fact that Φ is democratic ensures that this strategy will be effective [123].

Theorem 8.2 further guarantees graceful degradation due to loss of samples. Specifically, the theorem implies that reconstruction from any subset of CS measurements is stable to the loss of a potentially larger number of measurements than anticipated. To see this, suppose that an $M \times N$ matrix Φ is $(M - D, K, \delta)$-democratic, but consider the situation where $D + \widetilde{D}$ measurements are deleted. It is clear from the proof of Theorem 8.2 that if $\widetilde{D} < K$, then the resulting matrix Φ^{Γ} will satisfy the RIP of order $K - \widetilde{D}$ with constant δ. Thus, if we define $\widetilde{K} = (K - \widetilde{D})/2$, then as an example we have that from Theorem 3.2 the reconstruction error of BPDN in this setting is then bounded by

$$\|x - \widehat{x}\|_2 \leq C_3 \frac{\|x - x_{\widetilde{K}}\|_1}{\sqrt{\widetilde{K}}}, \tag{8.18}$$

where C_3 is an absolute constant depending on Φ that can be bounded using the constants derived in Theorem 8.2. Thus, if $\widetilde{D}$ is small then the additional error caused by deleting too many measurements will also be relatively small. To our knowledge, there is simply no analog to this kind of graceful degradation result for uniform sampling with linear reconstruction. When the number of deleted samples

exceeds D, there are no guarantees as to the accuracy of the reconstruction.

8.4.4 Simulations

As discussed previously, the democracy property is a stronger condition than the RIP. To demonstrate this, we perform a numerical simulation which illustrates this point. Specifically, we would like to compare the case where the measurements are deleted at random versus the case where the deleted measurements are selected by an adversary. Ideally, we would like to know whether the resulting matrices satisfy the RIP. Of course, this experiment is impossible to perform for two reasons: first, determining if a matrix satisfies the RIP is computationally intractable, as it would require checking all possible K-dimensional sub-matrices of Φ^Γ. Moreover, in the adversarial setting one would also have to search for the worst possible Γ as well, which is impossible for the same reason. Thus, we instead perform a far simpler experiment, which serves as a very rough proxy to the experiment we would like to perform.

The experiment proceeds over 100 trials as follows. We fix the parameters $N = 2048$ and $K = 13$ and vary M in the range $(0, 380)$. In each trial we draw a new matrix Φ with $\phi_{ij} \sim \mathcal{N}(0, 1/M)$ and a new signal with K nonzero coefficients, also drawn from a Gaussian distribution, and then the signal is normalized $\|x\|_2 = 1$. Over each set of trials we estimate two quantities:

1. the maximum D such that we achieve exact reconstruction for a randomly selected $(M - D) \times N$ submatrix of Φ on each of the 100 trials;

2. the maximum D such that we achieve exact reconstruction for $R = 300$ randomly selected $(M - D) \times N$ submatrices of Φ on each of the 100 trials..

Ideally, the second case should consider *all* $(M - D) \times N$ submatrices of Φ rather than just 300 submatrices, but as this is not possible (for reasons discussed above) we

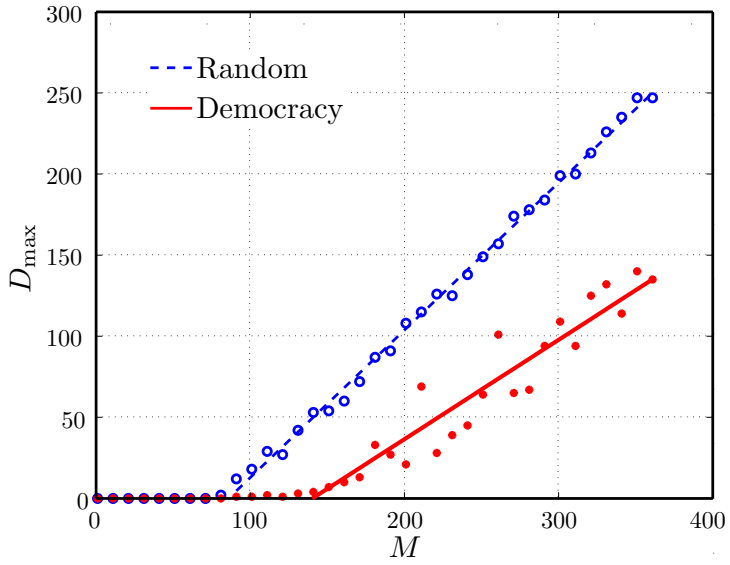

Figure 8.5: Maximum number of measurements that can be deleted $D_{\max}$ vs. number of measurements M for (a) exact recovery of one $(M - D) \times N$ submatrix of Φ and (b) exact recovery of $R = 300$ $(M - D) \times N$ submatrices of Φ.

simply perform a random sampling of the space of possible submatrices. Note also that exact recovery on one signal is also *not* proof that the matrix satisfies the RIP, although failure *is* proof that the matrix does not.

The results of this experiment are depicted in Figure 8.5. The circles denote data points with the empty circles corresponding to the random selection experiment and the solid circles corresponding to the democracy experiment. The lines denote the best linear fit for each data set where $D > 0$.

The maximum D corresponding to the random selection experiment grows linearly in M (with coefficient 1) once the minimum number of measurements required for RIP, denoted by M', is reached. This is because beyond this point at most $D = M - M'$ measurements can be discarded. As demonstrated by the plot, $M' \approx 90$ for this experiment. For the democracy experiment $M' \approx 150$, larger than for the RIP experiment. Furthermore, the maximum D for democracy grows more slowly than for the random selection case, which indicates that to be robust to the loss of *any* D measurements, CD additional measurements, with $C > 1$, are actually *necessary*.

Part IV

Sparse Signal Processing

Chapter 9

Compressive Detection, Classification, and Estimation

Despite the intense focus of the CS community on the problem of signal recovery, it is not actually necessary in many signal processing applications. In fact, most of DSP is actually concerned with solving *inference* problems, i.e., extracting only certain information from measurements. For example, we might aim to detect the presence of a signal of interest, classify among a set of possible candidate signals, estimate some function of the signal, or filter out a signal that is not of interest before further processing. While one could always attempt to recover the full signal from the compressive measurements and then solve such problems using traditional DSP techniques, this approach is typically suboptimal in terms of both accuracy and efficiency.

This thesis takes some initial steps towards a general framework for what we call *compressive signal processing* (CSP), an alternative approach in which signal processing problems are solved directly in the compressive measurement domain *without* first resorting to a full-scale signal reconstruction. We begin in this chapter[1] with

[1]This work was done in collaboration with Richard G. Baraniuk, Petros T. Boufounos, and Michael B. Wakin [133, 134].

an analysis of three fundamental signal processing problems: detection, classification, and estimation. In the case of signal detection and classification from random measurements in the presence of Gaussian noise, we derive the optimal detector/classifier and analyze its performance. We show that in the high SNR regime we can reliably detect/classify with far fewer measurements than are required for recovery. We also propose a simple and efficient approach to the estimation of linear functions of the signal from random measurements. We argue that in all of these settings, we can exploit sparsity and random measurements to enable the design of efficient, universal acquisition hardware. While these choices do not exhaust the set of canonical signal processing operations, we believe that they provide a strong initial foundation for CSP.

9.1 Compressive Signal Processing

9.1.1 Motivation

In what settings is it actually beneficial to take randomized, compressive measurements of a signal in order to solve an inference problem? One may argue that prior knowledge of the signal to be acquired or of the inference task to be solved could lead to a customized sensing protocol that very efficiently acquires the relevant information. For example, suppose we wish to acquire a signal $x \in \Sigma_K$ or $x \in \Psi(\Sigma_K)$ for some known basis Ψ. If we knew in advance which elements were nonzero, then the most efficient and direct measurement scheme would simply project the signal into the appropriate K-dimensional subspace. As a second example, suppose we wish to detect a known signal. If we knew in advance the signal template, then the optimal and most efficient measurement scheme would simply involve a receiving filter explicitly matched to the candidate signal.

Clearly, in cases where strong *a priori* information is available, customized sensing

protocols may be appropriate. However, a key objective of this thesis is to illustrate the *agnostic* and *universal* nature of random compressive measurements as a compact signal representation. These features enable the design of exceptionally *efficient* and *flexible* compressive sensing hardware that can be used for the acquisition of a variety of signal classes and applied to a variety of inference tasks.

As has been demonstrated in Part II, random measurements can be used to acquire any sparse signal without requiring advance knowledge of the locations of the nonzero coefficients. Thus, compressive measurements are agnostic in the sense that they capture the relevant information for the entire class Σ_K. We extend this concept to the CSP framework and demonstrate that it is possible to design agnostic measurement schemes that preserve the necessary structure of large signal classes in a variety of signal processing settings.

Furthermore, we observe that one can select a randomized measurement scheme without any prior knowledge of the signal class. For instance, in conventional CS it is not necessary to know the transform basis in which the signal has a sparse representation when acquiring the measurements. The only dependence is between the complexity of the signal class (e.g., the sparsity level of the signal) and the number of random measurements that must be acquired. Thus, random compressive measurements are universal in the sense that if one designs a measurement scheme at random, then with high probability it will preserve the structure of the signal class of interest, and thus explicit *a priori* knowledge of the signal class is unnecessary. We broaden this result and demonstrate that random measurements can universally capture the information relevant for many CSP applications without any prior knowledge of either the signal class or the ultimate signal processing task. In such cases, the requisite number of measurements scales efficiently with both the complexity of the signal and the complexity of the task to be performed.

It follows that, in contrast to the task-specific hardware used in many classical

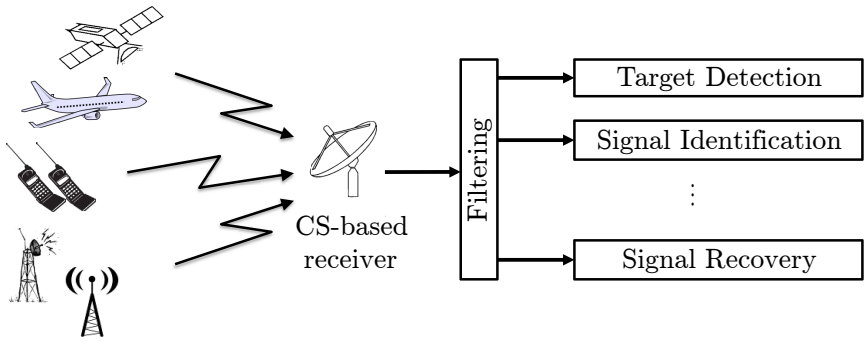

Figure 9.1: Example CSP application: Wideband signal monitoring.

acquisition systems, hardware designed to use a compressive measurement protocol can be extremely flexible. Returning to the binary detection scenario, for example, suppose that the signal template is unknown at the time of acquisition, or that one has a large number of candidate templates. Then what information should be collected at the sensor? A complete set of Nyquist samples would suffice, or a bank of matched filters could be employed. From a CSP standpoint, however, the solution is more elegant: one need only collect a small number of compressive measurements from which many candidate signals can be tested, many signal models can be posited, and many other inference tasks can be solved. What one loses in performance compared to a tailor-made matched filter, one may gain in simplicity and in the ability to adapt to future information about the problem at hand. In this sense, CSP impacts sensors in a similar manner as DSP impacted analog signal processing: expensive and inflexible analog components can be replaced by a universal, flexible, and programmable digital system.

9.1.2 Stylized application: Wideband signal monitoring

A stylized application to demonstrate the potential and applicability of the CSP framework is summarized in Figure 9.1. The figure schematically presents a wideband signal monitoring and processing system that receives signals from a variety of sources, including various television, radio, and cell-phone transmissions, radar signals, and

satellite communication signals. The extremely wide bandwidth monitored by such a system makes CS a natural approach for efficient signal acquisition [55].

In many cases, the system user might only be interested in extracting very small amounts of information from each signal. This can be efficiently performed using the tools we describe in the subsequent sections. For example, the user might be interested in detecting and classifying some of the signal sources, and in estimating some parameters, such as the location, of others. Full-scale signal recovery might be required for only a few of the signals in the monitored bandwidth. The detection, classification, and estimation tools developed below enable the system to perform these tasks much more efficiently in the compressive domain.

9.1.3 Context

There have been a number of related thrusts involving detection and classification using random measurements in a variety of settings. For example, in [61] sparsity is leveraged to perform classification with very few random measurements, while in [135, 136] random measurements are exploited to perform manifold-based image classification. In [124], small numbers of random measurements have also been noted as capturing sufficient information to allow robust face recognition. However, the most directly relevant work has been the discussions of detection in [137] and classification in [138]. We will contrast our results to those of [137, 138] below.

In this chapter we consider a variety of estimation and decision tasks. The data streaming community, which is concerned with efficient algorithms for processing large streams of data, has examined many similar problems over the past several years. The main differences with our work include: *(i)* data stream algorithms are typically designed to operate in noise-free environments on man-made digital signals, whereas we view compressive measurements as a sensing scheme that will operate in an inherently noisy environment; *(ii)* data stream algorithms typically provide

probabilistic guarantees, while we focus on providing deterministic guarantees; and *(iii)* data stream algorithms tend to tailor the measurement scheme to the task at hand, while we demonstrate that it is often possible to use the same measurements for a variety of signal processing tasks.

Finally, we note that in the remainder of this chapter, we will use the notation introduced in Section 4.4. This will allow us to state our results in a general manner that includes the sparse signal model, but also other signal models as described in Section 4.4.

9.2 Detection with Compressive Measurements

9.2.1 Problem setup and applications

We begin by examining the simplest of detection problems. We aim to distinguish between two hypotheses:

$$\mathcal{H}_0 : y = \Phi n$$
$$\mathcal{H}_1 : y = \Phi(s + n)$$

where $s \in \mathbb{R}^N$ is a known signal, $n \sim \mathcal{N}(0, \sigma^2 I_N)$ is i.i.d. Gaussian noise, and Φ is a known (fixed) measurement matrix. If s is known at the time of the design of Φ, then it is easy to show that the optimal design would be to set $\Phi = s^T$, which is just the *matched filter*. However, as mentioned in Section 9.1, we are often interested in universal or agnostic Φ. As an example, if we design hardware to implement the matched filter for a particular s, then we are very limited in what other signal processing tasks that hardware can perform. Even if we are only interested in detection, it is still possible that the signal s that we wish to detect may evolve over time. Thus, we will consider instead the case where Φ is designed without knowledge of s but is instead a

random matrix. From the results of Section 4.4, this will imply performance bounds that depend on how many measurements are acquired and the class $\mathcal{S}$ of possible s that we wish to detect.

9.2.2 Theory

To set notation, let

$$\begin{aligned}\mathbb{P}_F =& \mathbb{P}(\mathcal{H}_1 \text{ chosen when } \mathcal{H}_0 \text{ true}) \quad \text{and} \\ \mathbb{P}_D =& \mathbb{P}(\mathcal{H}_1 \text{ chosen when } \mathcal{H}_1 \text{ true})\end{aligned}$$

denote the *false alarm rate* and the *detection rate*, respectively. The *Neyman-Pearson* (NP) detector is the decision rule that maximizes $\mathbb{P}_D$ subject to the constraint that $\mathbb{P}_F \leq \alpha$. In order to derive the NP detector, we first observe that for our hypotheses, $\mathcal{H}_0$ and $\mathcal{H}_1$, we have the probability density functions[2]

$$f_0(y) = \frac{\exp\left(-\frac{1}{2}y^T(\sigma^2\Phi\Phi^T)^{-1}y\right)}{\det\left(\sigma^2\Phi\Phi^T\right)^{\frac{1}{2}}(2\pi)^{\frac{M}{2}}}$$

and

$$f_1(y) = \frac{\exp\left(-\frac{1}{2}(y-\Phi s)^T(\sigma^2\Phi\Phi^T)^{-1}(y-\Phi s)\right)}{\det\left(\sigma^2\Phi\Phi^T\right)^{\frac{1}{2}}(2\pi)^{\frac{M}{2}}},$$

where det denotes the matrix determinant. It is easy to show (see [139, 140], for example) that the NP-optimal decision rule is to compare the ratio $f_1(y)/f_0(y)$ to a threshold η, i.e, the *likelihood ratio test*:

$$\Lambda(y) = \frac{f_1(y)}{f_0(y)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \eta$$

[2]This formulation assumes that $\text{rank}(\Phi) = M$ so that $\Phi\Phi^T$ is invertible. If the entries of Φ are generated according to a continuous distribution and $M < N$, then this will be true with probability 1. This will also be true with high probability for discrete distributions provided that $M \ll N$. In the event that Φ is not full rank, appropriate adjustments can be made.

where η is chosen such that

$$\mathbb{P}_F = \int_{\Lambda(y)>\eta} f_0(y)\,dy = \alpha.$$

By taking a logarithm we obtain an equivalent test that simplifies to

$$y^T(\Phi\Phi^T)^{-1}\Phi s \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \sigma^2\log(\eta) + \frac{1}{2}s^T\Phi^T(\Phi\Phi^T)^{-1}\Phi s := \gamma.$$

We now define the compressive detector:

$$t := y^T(\Phi\Phi^T)^{-1}\Phi s. \tag{9.1}$$

It can be shown that t is a *sufficient statistic* for our detection problem, and thus t contains all of the information relevant for distinguishing between $\mathcal{H}_0$ and $\mathcal{H}_1$.

We must now set γ to achieve the desired performance. To simplify notation, we define

$$P_{\Phi^T} = \Phi^T(\Phi\Phi^T)^{-1}\Phi$$

as the orthogonal projection operator onto $\mathcal{R}(\Phi^T)$, i.e., the row space of Φ. Since $P_{\Phi^T} = P_{\Phi^T}^T$ and $P_{\Phi^T}^2 = P_{\Phi^T}$, we then have that

$$s^T\Phi^T(\Phi\Phi^T)^{-1}\Phi s = \|P_{\Phi^T}s\|_2^2. \tag{9.2}$$

Using this notation, it is easy to show that

$$t \sim \begin{cases} \mathcal{N}(0, \sigma^2\|P_{\Phi^T}s\|_2^2) & \text{under } \mathcal{H}_0 \\ \mathcal{N}(\|P_{\Phi^T}s\|_2^2, \sigma^2\|P_{\Phi^T}s\|_2^2) & \text{under } \mathcal{H}_1. \end{cases}$$

Thus we have

$$\begin{aligned}\mathbb{P}_F =& P(t > \gamma|\mathcal{H}_0) = Q\left(\frac{\gamma}{\sigma\|P_{\Phi^T}s\|_2}\right)\\ \mathbb{P}_D =& P(t > \gamma|\mathcal{H}_1) = Q\left(\frac{\gamma - \|P_{\Phi^T}s\|_2^2}{\sigma\|P_{\Phi^T}s\|_2}\right)\end{aligned}$$

where

$$Q(z) = \frac{1}{\sqrt{2\pi}}\int_z^\infty \exp\left(-u^2/2\right)\, du.$$

To determine the threshold, we set $\mathbb{P}_F = \alpha$, and thus

$$\gamma = \sigma\|P_{\Phi^T}s\|_2 Q^{-1}(\alpha),$$

resulting in

$$\mathbb{P}_D(\alpha) = Q\left(Q^{-1}(\alpha) - \|P_{\Phi^T}s\|_2/\sigma\right). \tag{9.3}$$

In general, this performance could be either quite good or quite poor depending on Φ. In particular, the larger $\|P_{\Phi^T}s\|_2$ is, then the better the performance. Recalling that P_{Φ^T} is the orthogonal projection onto the row space of Φ, we see that $\|P_{\Phi^T}s\|_2$ is simply the norm of the component of s that lies in the row space of Φ. This quantity is clearly at most $\|s\|_2$, which would yield the same performance as the traditional matched filter, but it could also be 0 if s lies in the null space of Φ. As we will see below, however, in the case where Φ is random, we can expect that $\|P_{\Phi^T}s\|_2$ concentrates around $\sqrt{M/N}\|s\|_2$.

Let us now define

$$\text{SNR} := \|s\|_2^2/\sigma^2. \tag{9.4}$$

We can bound the performance of the compressive detector as follows.

Theorem 9.1. *Suppose that* $\sqrt{N/M}P_{\Phi^T}$ *provides a* δ*-stable embedding of* $(\mathcal{S}, \{0\})$*. Then for any* $s \in \mathcal{S}$*, we can detect* s *with error rate*

$$\mathbb{P}_D(\alpha) \leq Q\left(Q^{-1}(\alpha) - \sqrt{1+\delta}\sqrt{\frac{M}{N}}\sqrt{SNR}\right) \tag{9.5}$$

and

$$\mathbb{P}_D(\alpha) \geq Q\left(Q^{-1}(\alpha) - \sqrt{1-\delta}\sqrt{\frac{M}{N}}\sqrt{SNR}\right). \tag{9.6}$$

Proof. By our assumption that $\sqrt{N/M}P_{\Phi^T}$ provides a δ-stable embedding of $(\mathcal{S}, \{0\})$, we know from (4.26) that

$$\sqrt{1-\delta}\|s\|_2 \leq \sqrt{\frac{N}{M}}\|P_{\Phi^T}s\|_2 \leq \sqrt{1+\delta}\|s\|_2. \tag{9.7}$$

Combining (9.7) with (9.3) and recalling the definition of the SNR from (9.4), the result follows. □

Theorem 9.1 tells us in a precise way how much information we lose by using random projections rather than the signal samples themselves, not in terms of our ability to recover the signal, but in terms of our ability to solve a detection problem. Specifically, for typical values of δ,

$$\mathbb{P}_D(\alpha) \approx Q\left(Q^{-1}(\alpha) - \sqrt{M/N}\sqrt{\mathrm{SNR}}\right), \tag{9.8}$$

which increases the miss probability by an amount determined by the SNR and the ratio M/N. Note that this is essentially the same phenomenon described in Chapter 7 — within the M-dimensional measurement subspace (as mapped to by $\sqrt{N/M}P_{\Phi^T}$), we will preserve the norms of the elements in $\mathcal{S}$. Meanwhile, the variance of the additive noise in this subspace is increased by a factor of N/M. Thus, our SNR decreases by a factor of M/N. In this case, however, there is a subtle difference in

that the impact of this decrease in the SNR has a nonlinear effect on $\mathbb{P}_D$ since it is passed through the Q function. Thus, in the high SNR regime it is possible to have $M \ll N$ while observing only a mild impact on the resulting $\mathbb{P}_D$.

In order to more clearly illustrate the behavior of $\mathbb{P}_D(\alpha)$ as a function of M, we also establish the following corollary of Theorem 9.1.

Corollary 9.1. *Suppose that $\sqrt{N/M}P_{\Phi^T}$ provides a δ-stable embedding of $(\mathcal{S}, \{0\})$. Then for any $s \in \mathcal{S}$, we can detect s with success rate*

$$\mathbb{P}_D(\alpha) \geq 1 - C_2 e^{-C_1 M/N}, \tag{9.9}$$

where C_1 and C_2 are constants depending only on α, δ, and the SNR.

Proof. We begin with the following bound from (13.48) of [141]

$$Q(z) \leq \frac{e^{-z^2/2}}{2}, \tag{9.10}$$

which allows us to bound $\mathbb{P}_D$ as follows. Let $C_1 = (1-\delta)\text{SNR}/2$. Then

$$\begin{aligned}
\mathbb{P}_D(\alpha) &\geq Q\left(Q^{-1}(\alpha) - \sqrt{2C_1 M/N}\right) \\
&= 1 - Q\left(\sqrt{2C_1 M/N} - Q^{-1}(\alpha)\right) \\
&\geq 1 - \frac{1}{2} e^{-C_1 M/N - \sqrt{2C_1 M/N}Q^{-1}(\alpha) + (Q^{-1}(\alpha))^2/2} \\
&\geq 1 - \frac{1}{2} e^{-C_1 M/N - \sqrt{2C_1}Q^{-1}(\alpha) + (Q^{-1}(\alpha))^2/2}.
\end{aligned}$$

Thus, if we let

$$C_2 = \frac{1}{2} e^{-Q^{-1}(\alpha)\left(Q^{-1}(\alpha)/2 - \sqrt{2C_1}\right)}, \tag{9.11}$$

we obtain the desired result. □

Thus, for a fixed SNR and signal length, the detection probability approaches 1

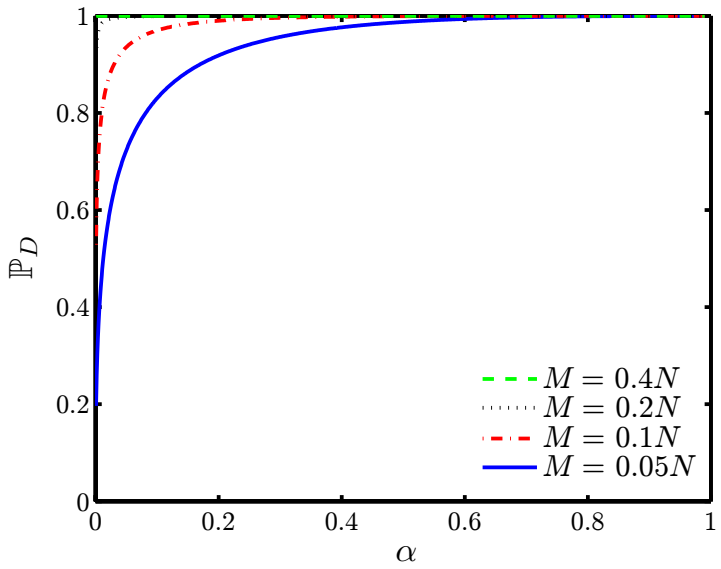

Figure 9.2: Effect of M on $\mathbb{P}_D(\alpha)$ predicted by (9.8) (SNR = 20dB).

exponentially fast as we increase the number of measurements.

9.2.3 Simulations and discussion

We first explore how M affects the performance of the compressive detector. As described above, decreasing M does cause a degradation in performance. However, as illustrated in Figure 9.2, in certain cases (relatively high SNR; 20 dB in this example) the compressive detector can perform almost as well as the traditional detector with a very small fraction of the number of measurements required by traditional detection. Specifically, in Figure 9.2 we illustrate the *receiver operating characteristic* (ROC) curve, i.e., the relationship between $\mathbb{P}_D$ and $\mathbb{P}_F$ predicted by (9.8). Observe that as M increases, the ROC curve approaches the upper-left corner, meaning that we can achieve very high detection rates while simultaneously keeping the false alarm rate very low. As M grows we see that we rapidly reach a regime where any additional increase in M yields only marginal improvements in the tradeoff between $\mathbb{P}_D$ and $\mathbb{P}_F$.

Furthermore, the exponential increase in the detection probability as we take more measurements is illustrated in Figure 9.3, which plots the performance predicted by

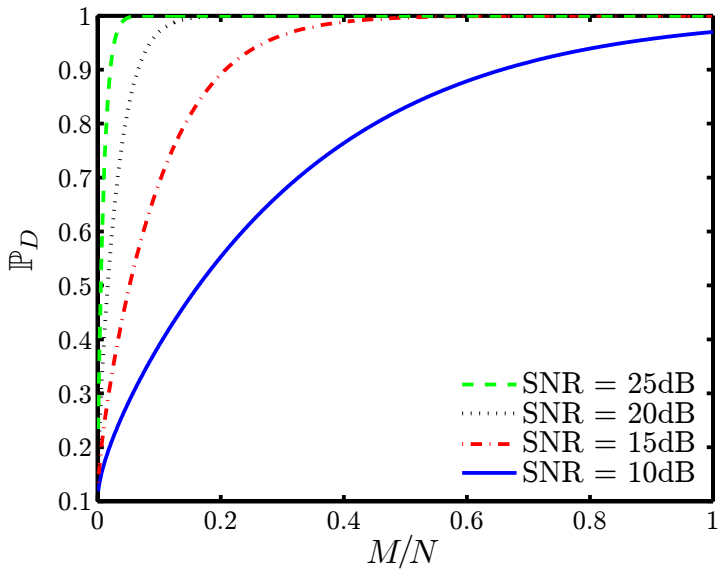

Figure 9.3: Effect of M on $\mathbb{P}_D$ predicted by (9.8) at several different SNR levels ($\alpha = 0.1$).

(9.8) for a range of SNRs with $\alpha = 0.1$. However, we again note that in practice this rate can be significantly affected by the SNR, which determines the constants in the bound of (9.9). These results are consistent with those obtained in [137], which also established that $\mathbb{P}_D$ should approach 1 exponentially fast as M is increased.

Finally, we close by noting that for any given instance of Φ, its ROC curve may be better or worse than that predicted by (9.8). However, with high probability it is tightly concentrated around the expected performance curve. Figure 9.4 illustrates this for the case where s is fixed, the SNR is 20dB, Φ has i.i.d. Gaussian entries, $M = 0.05N$, and $N = 1000$. The predicted ROC curve is illustrated along with curves displaying the best and worst ROC curves obtained over 100 independent draws of Φ. We see that our performance is never significantly different from what we expect. Furthermore, we have also observed that these bounds grow significantly tighter as we increase N; so for large problems the difference between the predicted and actual curves will be insignificant.

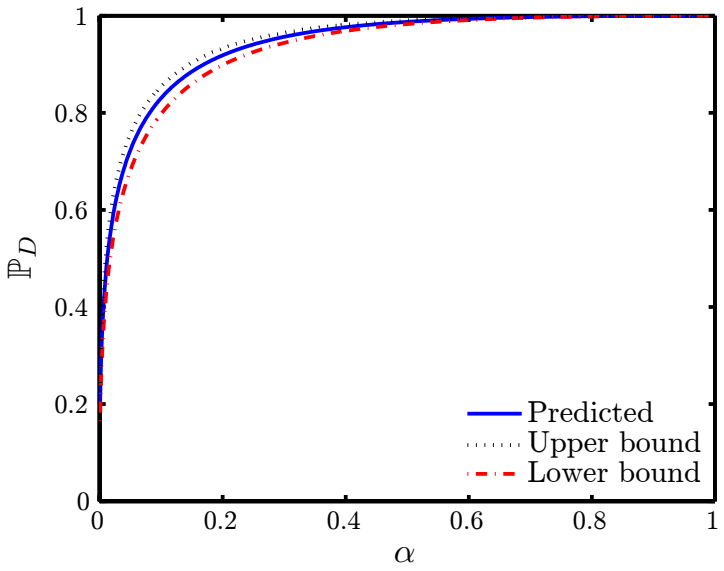

Figure 9.4: Concentration of ROC curves for random Φ near the expected ROC curve (SNR = 20dB, $M = 0.05N$, $N = 1000$).

9.3 Classification with Compressive Measurements

9.3.1 Problem setup and applications

We can easily generalize the setting of Section 9.2 to the problem of binary classification. Specifically, if we wish to distinguish between $\Phi(s_0+n)$ and $\Phi(s_1+n)$, then it is equivalent to be able to distinguish $\Phi(s_0+n)-\Phi s_0 = \Phi n$ and $\Phi(s_1-s_0+n)$. Thus, the conclusions for the case of binary classification are identical to those discussed in Section 9.2.

More generally, suppose that we would like to distinguish between the hypotheses:

$$\widetilde{\mathcal{H}}_i \ : \ y = \Phi(s_i + n),$$

for $i = 1, 2, \ldots, R$, where each $s_i \in \mathcal{S}$ is one of our known signals and as before, $n \sim \mathcal{N}(0, \sigma^2 I_N)$ is i.i.d. Gaussian noise and Φ is a known $M \times N$ matrix.

It is straightforward to show (see [139, 140], for example), in the case where each hypothesis is equally likely, that the classifier with minimum probability of error

selects the $\widetilde{\mathcal{H}_i}$ that minimizes

$$t_i := (y - \Phi s_i)^T (\Phi\Phi^T)^{-1} (y - \Phi s_i). \tag{9.12}$$

If the rows of Φ are orthogonal and have equal norm, then this reduces to identifying which Φs_i is closest to y. The $(\Phi\Phi^T)^{-1}$ term arises when the rows of Φ are not orthogonal because the noise is no longer uncorrelated.

As an alternative illustration of the classifier behavior, let us suppose that $y = \Phi x$ for some $x \in \mathbb{R}^N$. Then, starting with (9.12), we have

$$\begin{aligned} t_i &= (y - \Phi s_i)^T (\Phi\Phi^T)^{-1} (y - \Phi s_i) \\ &= (\Phi x - \Phi s_i)^T (\Phi\Phi^T)^{-1} (\Phi x - \Phi s_i) \\ &= (x - s_i)^T \Phi^T (\Phi\Phi^T)^{-1} \Phi (x - s_i) \\ &= \|P_{\Phi^T} x - P_{\Phi^T} s_i\|_2^2, \end{aligned} \tag{9.13}$$

where (9.13) follows from the same argument as (9.2). Thus, we can equivalently think of the classifier as simply projecting x and each candidate signal s_i onto the row space of Φ and then classifying according to the nearest neighbor in this space.

9.3.2 Theory

While in general it is difficult to find analytical expressions for the probability of error even in non-compressive classification settings, we can provide a bound for the performance of the compressive classifier as follows.

Theorem 9.2. *Suppose that $\sqrt{N/M} P_{\Phi^T}$ provides a δ-stable embedding of $(\mathcal{S}, \mathcal{S})$, and let $R = |\mathcal{S}|$. Let*

$$d = \min_{i,j} \|s_i - s_j\|_2 \tag{9.14}$$

denote the minimum separation among the s_i. For some $i^* \in \{1, 2, \ldots, R\}$, let $y =$

$\Phi(s_{i^*} + n)$, *where* $n \sim \mathcal{N}(0, \sigma^2 I_N)$ *is i.i.d. Gaussian noise. Then with probability at least*

$$1 - \left(\frac{R-1}{2}\right) e^{-d^2(1-\delta)M/(8\sigma^2 N)}, \tag{9.15}$$

the signal can be correctly classified, i.e.,

$$i^* = \underset{i \in \{1,2,\ldots,R\}}{\arg\min}\ t_i. \tag{9.16}$$

Proof. Let $j \neq i^*$. We will argue that $t_j > t_{i^*}$ with high probability. From (9.13) we have that

$$t_{i^*} = \|P_{\Phi^T} n\|_2^2$$

and

$$\begin{aligned} t_j &= \|P_{\Phi^T}(s_{i^*} - s_j + n)\|_2^2 \\ &= \|P_{\Phi^T}(s_{i^*} - s_j) + P_{\Phi^T} n\|_2^2 \\ &= \|\tau + P_{\Phi^T} n\|_2^2, \end{aligned}$$

where we have defined $\tau = P_{\Phi^T}(s_{i^*} - s_j)$ to simplify notation. Let us define $P_\tau = \tau\tau^T/\|\tau\|_2^2$ as the orthogonal projection onto the 1-dimensional span of τ, and $P_{\tau^\perp} = (I_N - P_\tau)$. Then we have

$$t_{i^*} = \|P_\tau P_{\Phi^T} n\|_2^2 + \|P_{\tau^\perp} P_{\Phi^T} n\|_2^2$$

and

$$\begin{aligned} t_j &= \|P_\tau(\tau + P_{\Phi^T} n)\|_2^2 + \|P_{\tau^\perp}(\tau + P_{\Phi^T} n)\|_2^2 \\ &= \|\tau + P_\tau P_{\Phi^T} n\|_2^2 + \|P_{\tau^\perp} P_{\Phi^T} n\|_2^2. \end{aligned}$$

Thus, $t_j \leq t_{i^*}$ if and only if

$$\|\tau + P_\tau P_{\Phi^T} n\|_2^2 \leq \|P_\tau P_{\Phi^T} n\|_2^2,$$

or equivalently, if

$$\left\| \frac{\tau^T}{\|\tau\|_2}(\tau + P_\tau P_{\Phi^T} n) \right\|_2^2 \leq \left\| \frac{\tau^T}{\|\tau\|_2} P_\tau P_{\Phi^T} n \right\|_2^2,$$

or equivalently, if

$$\left| \|\tau\|_2 + \frac{\tau^T}{\|\tau\|_2} P_{\Phi^T} n \right| \leq \left| \frac{\tau^T}{\|\tau\|_2} P_{\Phi^T} n \right|,$$

or equivalently, if

$$\frac{\tau^T}{\|\tau\|_2} P_{\Phi^T} n \leq -\frac{\|\tau\|_2}{2}.$$

The quantity $\frac{\tau^T}{\|\tau\|_2} P_{\Phi^T} n$ is a scalar, zero-mean Gaussian random variable with variance

$$\frac{\tau^T}{\|\tau\|_2} P_{\Phi^T} (\sigma^2 I_N) P_{\Phi^T}^T \frac{\tau}{\|\tau\|_2} = \frac{\tau^T P_{\Phi^T} \tau \sigma^2}{\|\tau\|_2^2} = \sigma^2.$$

Because $\sqrt{N/M} P_{\Phi^T}$ provides a δ-stable embedding of $(\mathcal{S}, \mathcal{S})$, and by our assumption that $\|s_{i^*} - s_j\|_2 \geq d$, we have that $\|\tau\|_2^2 \geq d^2(1-\delta)M/N$. Thus, using also (9.10), we have

$$\begin{aligned}
\mathbb{P}(t_j \leq t_{i^*}) &= \mathbb{P}\left(\frac{\tau^T}{\|\tau\|_2} P_{\Phi^T} n \leq -\frac{\|\tau\|_2}{2} \right) \\
&= Q\left(\frac{\|\tau\|_2}{2\sigma} \right) \\
&\leq \frac{1}{2} e^{-\|\tau\|_2^2/(8\sigma^2)} \\
&\leq \frac{1}{2} e^{-d^2(1-\delta)M/(8\sigma^2 N)}.
\end{aligned}$$

Finally, because t_{i^*} is compared to $R-1$ other candidates, we use a union bound to conclude that (9.16) holds with probability exceeding that given in (9.15). $\square$

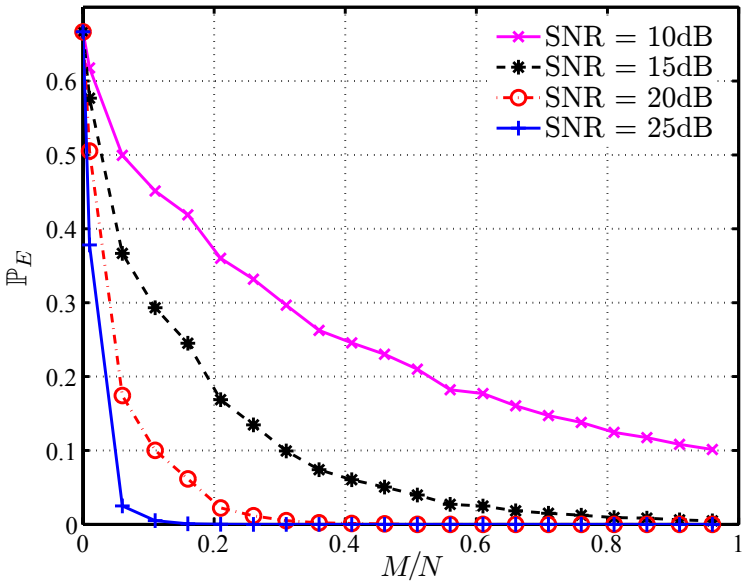

Figure 9.5: Effect of M on $\mathbb{P}_E$ (the probability of error of a compressive domain classifier) for $R = 3$ signals at several different SNR levels, where $\text{SNR} = 10\log_{10}(d^2/\sigma^2)$.

9.3.3 Simulations and discussion

In Figure 9.5 we display experimental results for classification among $R = 3$ test signals of length $N = 1000$. The signals s_1, s_2, and s_3 are drawn according to a Gaussian distribution with mean 0 and variance 1 and then fixed. For each value of M, a single Gaussian Φ is drawn and then the probability of error $\mathbb{P}_E$ is computed by averaging the results over 10^6 realizations of the noise vector n. The error rates are very similar in spirit to those for detection (see Figure 9.3). The results agree with Theorem 9.2, in which we demonstrate that, as was the case for detection, as M increases the probability of error decays exponentially fast. This also agrees with the related results of [138].

9.4 Estimation with Compressive Measurements

9.4.1 Problem setup and applications

While many signal processing problems can be reduced to a detection or classification problem, in some cases we cannot reduce our task to selecting among a finite set of hypotheses. Rather, we might be interested in *estimating* some function of the data. In this section we will focus on estimating a *linear* function of the data from compressive measurements.

Suppose that we observe $y = \Phi s$ and wish to estimate $\langle \ell, s \rangle$ from the measurements y, where $\ell \in \mathbb{R}^N$ is a fixed test vector. In the case where Φ is a random matrix, a natural estimator is essentially the same as the compressive detector. Specifically, suppose we have a set $\mathcal{L}$ of $|\mathcal{L}|$ linear functions we would like to estimate from y. Example applications include computing the coefficients of a basis or frame representation of the signal, estimating the signal energy in a particular linear subspace, parametric modeling, and so on. One potential estimator for this scenario, which is essentially a simple generalization of the compressive detector in (9.1), is given by

$$\frac{N}{M} y^T (\Phi \Phi^T)^{-1} \Phi \ell_i, \tag{9.17}$$

for $i = 1, 2, \ldots, |\mathcal{L}|$. While this approach, which we shall refer to as the *orthogonalized* estimator, has certain advantages, it is also enlightening to consider an even simpler estimator, given by

$$\langle y, \Phi \ell_i \rangle . \tag{9.18}$$

We shall refer to this approach as the *direct* estimator since it eliminates the orthogonalization step by directly correlating the compressive measurements with $\Phi \ell_i$. We will provide a more detailed experimental comparison of these two approaches below, but in the proof of Corollary 9.2 we focus only on the direct estimator.

9.4.2 Theory

We now provide bounds on the performance of our simple estimator.[3] This bound is a generalization of Lemma 6.1 to the setting of more general stable embeddings. The proof is omitted as it is essentially identical to that of Lemma 6.1.

Corollary 9.2. *Suppose that $\ell \in \mathcal{L}$ and $s \in \mathcal{S}$ and that Φ is a δ-stable embedding of $(\mathcal{L}, \mathcal{S} \cup -\mathcal{S})$. Then*

$$|\langle \Phi\ell, \Phi s \rangle - \langle \ell, s \rangle| \leq \delta \|\ell\|_2 \|s\|_2. \tag{9.19}$$

One way of interpreting our result is that the angle between two vectors can be estimated accurately; this is formalized as follows.

Corollary 9.3. *Suppose that $\ell \in \mathcal{L}$ and $s \in \mathcal{S}$ and that Φ is a δ-stable embedding of $(\mathcal{L} \cup \{0\}, \mathcal{S} \cup -\mathcal{S} \cup \{0\})$. Then*

$$|\cos \measuredangle(\Phi\ell, \Phi s) - \cos \measuredangle(\ell, s)| \leq 2\delta,$$

where $\measuredangle(\cdot, \cdot)$ denotes the angle between two vectors.

Proof. By definition, we have

$$\cos \measuredangle(\ell, s) = \frac{\langle \ell, s \rangle}{\|\ell\|_2 \|s\|_2}$$

and

$$\cos \measuredangle(\Phi\ell, \Phi s) = \frac{\langle \Phi\ell, \Phi s \rangle}{\|\Phi\ell\|_2 \|\Phi s\|_2}.$$

Thus, from (9.19) we have

$$\left| \frac{\langle \Phi\ell, \Phi s \rangle}{\|\ell\|_2 \|s\|_2} - \cos \measuredangle(\ell, s) \right| \leq \delta. \tag{9.20}$$

[3]Note that the same guarantee can be established for the orthogonalized estimator under the assumption that $\sqrt{N/M} P_{\Phi^T}$ is a δ-stable embedding of $(\mathcal{L}, \mathcal{S} \cup -\mathcal{S})$.

Now, using the fact that Φ is a δ-stable embedding, we can show that

$$\frac{(1-\delta)}{\|\Phi\ell\|_2\|\Phi s\|_2} \leq \frac{1}{\|\ell\|_2\|s\|_2} \leq \frac{(1+\delta)}{\|\Phi\ell\|_2\|\Phi s\|_2},$$

from which we infer that

$$\left| \frac{\langle \Phi\ell, \Phi s\rangle}{\|\ell\|_2\|s\|_2} - \frac{\langle \Phi\ell, \Phi s\rangle}{\|\Phi\ell\|_2\|\Phi s\|_2} \right| \leq \delta \frac{\langle \Phi\ell, \Phi s\rangle}{\|\Phi\ell\|_2\|\Phi s\|_2} \leq \delta. \tag{9.21}$$

Therefore, combining (9.20) and (9.21) using the triangle inequality, the desired result follows. □

While Corollary 9.2 suggests that the absolute error in estimating $\langle \ell, s\rangle$ must scale with $\|\ell\|_2\|s\|_2$, this is probably the best we can expect. If the $\|\ell\|_2\|s\|_2$ terms were omitted on the right hand side of (9.19), then one could estimate $\langle \ell, s\rangle$ with arbitrary accuracy using the following strategy: *(i)* choose a large positive constant C_{big}, *(ii)* estimate the inner product $\langle C_{\text{big}}\ell, C_{\text{big}}s\rangle$, obtaining an accuracy δ, and then *(iii)* divide the estimate by C_{big}^2 to estimate $\langle \ell, s\rangle$ with accuracy δ/C_{big}^2. Similarly, it is not possible to replace the right hand side of (9.19) with an expression proportional merely to $\langle \ell, s\rangle$, as this would imply that $\langle \Phi\ell, \Phi s\rangle = \langle \ell, s\rangle$ exactly when $\langle \ell, s\rangle = 0$, and unfortunately this is not the case. (Were this possible, one could exploit this fact to immediately identify the nonzero locations in a sparse signal by letting $\ell_i = e_i$, the i^{th} canonical basis vector, for $i = 1, 2, \ldots, N$.)

9.4.3 Simulations and discussion

In Figure 9.6 we display the average estimation error for the orthogonalized and direct estimators, i.e.,

$$\left|(N/M)s^T\Phi^T(\Phi\Phi^T)^{-1}\Phi\ell - \langle \ell, s\rangle\right| / \|s\|_2\|\ell\|_2$$

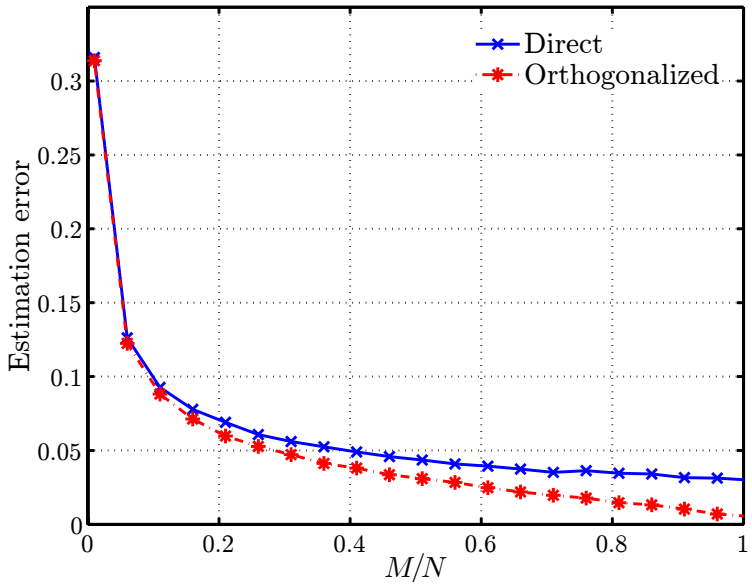

Figure 9.6: Average error in the estimate of the mean of a fixed signal s.

and

$$\left|\langle \Phi\ell, \Phi s\rangle - \langle \ell, s\rangle\right| / \|s\|_2 \|\ell\|_2$$

respectively. The signal s is a length $N = 1000$ vector with entries distributed according to a Gaussian distribution with mean 1 and unit variance. We choose $\ell = [\frac{1}{N} \ \frac{1}{N} \cdots \frac{1}{N}]^T$ to compute the mean of s. The result displayed is the mean error averaged over 10^4 different draws of Gaussian Φ with s fixed. Note that we obtain nearly identical results for other candidate ℓ, including ℓ both highly correlated with s and ℓ nearly orthogonal to s. In all cases, as M increases, the error decays because the random matrices Φ become δ-stable embeddings of $\{s\}$ for smaller values of δ. Note that for small values of M, there is very little difference between the orthogonalized and direct estimators. The orthogonalized estimator only provides notable improvement when M is large, in which case the computational difference is significant. In this case one must weigh the relative importance of speed versus accuracy in order to judge which approach is best, so the proper choice will ultimately be dependent on the application.

In the case where $|\mathcal{L}| = 1$, Corollary 9.2 is a deterministic version of Theorem 4.5

of [142] and Lemma 3.1 of [143], which both show that for certain random constructions of Φ, with probability at least $1 - \rho$,

$$|\langle \Phi \ell, \Phi s \rangle - \langle \ell, s \rangle| \leq \delta \|\ell\|_2 \|s\|_2. \tag{9.22}$$

In [142] $\rho = 2\delta^2/M$, while in [143] more sophisticated methods are used to achieve a bound on ρ of the form $\rho \leq 2e^{-cM\delta^2}$ as in (4.16). Our result extends these results to the wider class of sub-Gaussian matrices. Furthermore, our approach generalizes naturally to simultaneously estimating multiple linear functions of the data.

Specifically, it is straightforward to extend our analysis beyond the estimation of scalar-valued linear functions to more general linear operators. Any finite-dimensional linear operator on a signal $x \in \mathbb{R}^N$ can be represented as a matrix multiplication Lx, where L has size $Z \times N$ for some Z. Decomposing L in terms of its rows, this computation can be expressed as

$$Lx = \begin{bmatrix} \ell_1^T \\ \ell_2^T \\ \vdots \\ \ell_Z^T \end{bmatrix} x = \begin{bmatrix} \langle \ell_1, x \rangle \\ \langle \ell_2, x \rangle \\ \vdots \\ \langle \ell_Z, x \rangle \end{bmatrix}.$$

From this point, the bound (9.19) can be applied to each component of the resulting vector. It is also interesting to note that if $L = I$, then we can observe that

$$\|\Phi^T \Phi x - x\|_\infty \leq \delta \|x\|_2.$$

This could be used to establish deterministic bounds on the performance of the thresholding signal recovery algorithm described in [143], which simply thresholds $\Phi^T y$ to keep only the K largest elements. Moreover, we have already applied the essence of this result in Chapter 6 in our analysis of OMP.

Note that one could clearly consider more sophisticated estimators, even for the simple problem of linear function estimation. Specifically, in the case where $\mathcal{S} = \Sigma_K$, then one could obtain an exact estimate by first recovering the signal. In general, the techniques described in this section are highly efficient but do not necessarily fully exploit the structure in $\mathcal{S}$, which leaves significant room for improvement for specific choices of $\mathcal{S}$.

Chapter 10

Compressive Filtering

This chapter[1] analyzes the problem of filtering compressive measurements. We begin with a simple method for suppressing sparse interference. We demonstrate the relationship between this method and a key step in orthogonal greedy algorithms and illustrate its application to the problem of signal recovery in the presence of interference, or equivalently, signal recovery with partially known support. We then generalize this method to more general filtering methods, with a particular focus on the cancellation of bandlimited, but not necessarily sparse, interference. These filtering procedures ultimately facilitate the separation of signals after they have been acquired in the compressive domain so that each signal can be processed by the appropriate algorithm, depending on the information sought by the user.

10.1 Subspace Filtering

10.1.1 Problem setup and applications

In practice, it is often the case that the signal we wish to acquire is contaminated with interference. The universal nature of compressive measurements, while often

[1]This work was done in collaboration with Richard G. Baraniuk, Petros T. Boufounos, and Michael B. Wakin [113, 134].

advantageous, can also increase our susceptibility to interference and significantly affect the performance of algorithms such as those described in Sections 9.2–9.4. It is therefore desirable to remove unwanted signal components from the compressive measurements before they are processed further.

More formally, suppose that the signal $x \in \mathbb{R}^N$ consists of two components:

$$x = x_S + x_I,$$

where x_S represents the *signal of interest* and x_I represents an unwanted signal that we would like to reject. We refer to x_I as *interference* in the remainder of this section, although it might be the signal of interest for a different system module. Supposing we acquire measurements of both components simultaneously

$$y = \Phi(x_S + x_I), \tag{10.1}$$

our goal is to remove the contribution of x_I from the measurements y while preserving the information about x_S. In this section, we will assume that $x_S \in \mathcal{S}_S$ and that $x_I \in \mathcal{S}_I$. In our discussion, we will further assume that Φ is a δ-stable embedding of $(\widetilde{\mathcal{S}}_S, \mathcal{S}_I)$, where $\widetilde{\mathcal{S}}_S$ is a set with a simple relationship to $\mathcal{S}_S$ and $\mathcal{S}_I$.

While one could consider more general interference models, we restrict our attention to the case where either the interfering signal or the signal of interest lives in a known subspace. For example, suppose we have obtained measurements of a radio signal that has been corrupted by narrow band interference such as a TV or radio station operating at a known carrier frequency. In this case we can project the compressive measurements into a subspace orthogonal to the interference, and hence eliminate the contribution of the interference to the measurements. We further demonstrate that provided that the signal of interest is orthogonal to the set of possible interference signals, the projection operator maintains a stable embedding

for the set of signals of interest. Thus, the projected measurements retain sufficient information to enable the use of efficient compressive-domain algorithms for further processing.

10.1.2 Theory

We first consider the case where $\mathcal{S}_I$ is a K_I-dimensional subspace, and we place no restrictions on the set $\mathcal{S}_S$. We will later see that by symmetry the methods we develop for this case will have implications for the setting where $\mathcal{S}_S$ is a K_S-dimensional subspace and where $\mathcal{S}_I$ is a more general set.

We filter out the interference by constructing a linear operator P that operates on the measurements y. The design of P is based solely on the measurement matrix Φ and knowledge of the subspace $\mathcal{S}_I$. Our goal is to construct a P that maps Φx_I to zero for any $x_I \in \mathcal{S}_I$. To simplify notation, we assume that Ψ_I is an $N \times K_I$ matrix whose columns form an orthonormal basis for the K_I-dimensional subspace $\mathcal{S}_I$, and we define the $M \times K_I$ matrix $\Omega = \Phi\Psi_I$. We recall the definitions of

$$P_\Omega = \Omega\Omega^\dagger \tag{10.2}$$

and

$$P_{\Omega^\perp} = I - P_\Omega = I - \Omega\Omega^\dagger \tag{10.3}$$

as the orthogonal projection operators onto $\mathcal{R}(\Omega)$ and its orthogonal complement. The resulting $P_{\Omega^\perp}$ is our desired operator P: it is an orthogonal projection operator onto the orthogonal complement of $\mathcal{R}(\Omega)$, and its null space equals $\mathcal{R}(\Omega)$.

Using Corollary 9.2, we now show that the fact that Φ is a stable embedding allows us to argue that $P_{\Omega^\perp}$ preserves the structure of $\widetilde{\mathcal{S}}_S = P_{\mathcal{S}_I^\perp}\mathcal{S}_S$ (where $\mathcal{S}_I^\perp$ denotes the orthogonal complement of $\mathcal{S}_I$ and $P_{\mathcal{S}_I^\perp}$ denotes the orthogonal projection onto $\mathcal{S}_I^\perp$),

while simultaneously cancelling out signals from $\mathcal{S}_I$.[2] Additionally, P_Ω preserves the structure in $\mathcal{S}_I$ while nearly cancelling out signals from $\widetilde{\mathcal{S}}_S$.

Theorem 10.1. *Suppose that Φ is a δ-stable embedding of $(\widetilde{\mathcal{S}}_S \cup \{0\}, \mathcal{S}_I)$, where $\mathcal{S}_I$ is a K_I-dimensional subspace of $\mathbb{R}^N$ with orthonormal basis Ψ_I. Set $\Omega = \Phi\Psi_I$ and define P_Ω and $P_{\Omega^\perp}$ as in (10.2) and (10.3). For any $x \in \mathcal{S}_S \oplus \mathcal{S}_I$ we can write $x = \widetilde{x}_S + \widetilde{x}_I$, where $\widetilde{x}_S \in \widetilde{\mathcal{S}}_S$ and $\widetilde{x}_I \in \mathcal{S}_I$. Then*

$$P_{\Omega^\perp}\Phi x = P_{\Omega^\perp}\Phi\widetilde{x}_S \tag{10.4}$$

and

$$P_\Omega\Phi x = \Phi\widetilde{x}_I + P_\Omega\Phi\widetilde{x}_S. \tag{10.5}$$

Furthermore,

$$1 - \frac{\delta}{1-\delta} \leq \frac{\|P_{\Omega^\perp}\Phi\widetilde{x}_S\|_2^2}{\|\widetilde{x}_S\|_2^2} \leq 1 + \delta \tag{10.6}$$

and

$$\frac{\|P_\Omega\Phi\widetilde{x}_S\|_2^2}{\|\widetilde{x}_S\|_2^2} \leq \delta^2\frac{1+\delta}{(1-\delta)^2}. \tag{10.7}$$

Proof. We begin by observing that since $\widetilde{\mathcal{S}}_S$ and $\mathcal{S}_I$ are orthogonal, the decomposition $x = \widetilde{x}_S + \widetilde{x}_I$ is unique. Furthermore, since $\widetilde{x}_I \in \mathcal{S}_I$, we have that $\Phi\widetilde{x}_I \in \mathcal{R}(\Omega)$ and hence by the design of $P_{\Omega^\perp}$, $P_{\Omega^\perp}\Phi\widetilde{x}_I = 0$ and $P_\Omega\Phi\widetilde{x}_I = \Phi\widetilde{x}_I$, which establishes (10.4) and (10.5).

In order to establish (10.6) and (10.7), we decompose $\Phi\widetilde{x}_S$ as $\Phi\widetilde{x}_S = P_\Omega\Phi\widetilde{x}_S + P_{\Omega^\perp}\Phi\widetilde{x}_S$. Since P_Ω is an orthogonal projection we can write

$$\|\Phi\widetilde{x}_S\|_2^2 = \|P_\Omega\Phi\widetilde{x}_S\|_2^2 + \|P_{\Omega^\perp}\Phi\widetilde{x}_S\|_2^2. \tag{10.8}$$

[2]Note that we do not claim that $P_{\Omega^\perp}$ preserves the structure of $\mathcal{S}_S$, but rather the structure of $\widetilde{\mathcal{S}}_S$. This is because we do not restrict $\mathcal{S}_S$ to be orthogonal to the subspace $\mathcal{S}_I$ which we cancel. Clearly, we cannot preserve the structure of the component of $\mathcal{S}_S$ that lies within $\mathcal{S}_I$ while simultaneously eliminating interference from $\mathcal{S}_I$.

Furthermore, note that $P_\Omega^T = P_\Omega$ and $P_\Omega^2 = P_\Omega$, so that

$$\langle P_\Omega \Phi \widetilde{x}_S, \Phi \widetilde{x}_S \rangle = \|P_\Omega \Phi \widetilde{x}_S\|_2^2. \tag{10.9}$$

Since P_Ω is a projection onto $\mathcal{R}(\Omega)$ there exists a $z \in \mathcal{S}_I$ such that $P_\Omega \Phi \widetilde{x}_S = \Phi z$. Since $\widetilde{x}_S \in \widetilde{\mathcal{S}}_S$, we have that $\langle \widetilde{x}_S, z \rangle = 0$, and since $\mathcal{S}_I$ is a subspace, $\mathcal{S}_I = \mathcal{S}_I \cup -\mathcal{S}_I$, and so we may apply Corollary 9.2 to obtain

$$|\langle P_\Omega \Phi \widetilde{x}_S, \Phi \widetilde{x}_S \rangle| = |\langle \Phi z, \Phi \widetilde{x}_S \rangle| \leq \delta \|z\|_2 \|\widetilde{x}_S\|_2.$$

Since $0 \in \mathcal{S}_I$ and Φ is a δ-stable embedding of $(\widetilde{\mathcal{S}}_S \cup \{0\}, \mathcal{S}_I)$, we have that

$$\|z\|_2 \|\widetilde{x}_S\|_2 \leq \frac{\|\Phi z\|_2 \|\Phi \widetilde{x}_S\|_2}{1-\delta}.$$

Recalling that $\Phi z = P_\Omega \Phi \widetilde{x}_S$, we obtain

$$\frac{|\langle P_\Omega \Phi \widetilde{x}_S, \Phi \widetilde{x}_S \rangle|}{\|P_\Omega \Phi \widetilde{x}_S\|_2 \|\Phi \widetilde{x}_S\|_2} \leq \frac{\delta}{1-\delta}.$$

Combining this with (10.9), we obtain

$$\|P_\Omega \Phi \widetilde{x}_S\|_2 \leq \frac{\delta}{1-\delta} \|\Phi \widetilde{x}_S\|_2.$$

Since $\widetilde{x}_S \in \widetilde{\mathcal{S}}_S$, $\|\Phi \widetilde{x}_S\|_2 \leq \sqrt{1+\delta} \|\widetilde{x}_S\|_2$, and thus we obtain (10.7). Since we trivially have that $\|P_\Omega \Phi \widetilde{x}_S\|_2 \geq 0$, we can combine this with (10.8) to obtain

$$\left(1 - \left(\frac{\delta}{1-\delta}\right)^2\right) \|\Phi \widetilde{x}_S\|_2^2 \leq \|P_{\Omega^\perp} \Phi \widetilde{x}_S\|_2^2 \leq \|\Phi \widetilde{x}_S\|_2^2.$$

Again, since $\widetilde{x}_S \in \widetilde{\mathcal{S}}_S$, we have that

$$\left(1 - \left(\frac{\delta}{1-\delta}\right)^2\right)(1-\delta) \leq \frac{\|P_{\Omega^\perp}\Phi\widetilde{x}_S\|_2^2}{\|\widetilde{x}_S\|_2^2} \leq 1+\delta,$$

which simplifies to yield (10.6). □

Corollary 10.1. *Suppose that Φ is a δ-stable embedding of $(\widetilde{\mathcal{S}}_S \cup \{0\}, \mathcal{S}_I)$, where $\mathcal{S}_I$ is a K_I-dimensional subspace of $\mathbb{R}^N$ with orthonormal basis Ψ_I. Set $\Omega = \Phi\Psi_I$ and define P_Ω and $P_{\Omega^\perp}$ as in (10.2) and (10.3). Then $P_{\Omega^\perp}\Phi$ is a $\delta/(1-\delta)$-stable embedding of $(\widetilde{\mathcal{S}}_S, \{0\})$ and $P_\Omega\Phi$ is a δ-stable embedding of $(\mathcal{S}_I, \{0\})$.*

Proof. This follows from Theorem 10.1 by picking $x \in \widetilde{\mathcal{S}}_S$, in which case $x = \widetilde{x}_S$, or picking $x \in \mathcal{S}_I$, in which case $x = \widetilde{x}_I$. □

Theorem 10.1 and Corollary 10.1 have a number of practical benefits. For example, if we are interested in solving an inference problem based only on the signal x_S, then we can use P_Ω or $P_{\Omega^\perp}$ to filter out the interference and then apply the compressive domain inference techniques developed above. The performance of these techniques will be significantly improved by eliminating the interference due to x_I. Furthermore, this result also has implications for the problem of signal recovery, as demonstrated by the following corollary, which is a generalization of Lemma 6.2 (the two are equivalent in the case where Ψ_I is a submatrix of Φ).

Corollary 10.2. *Suppose that Ψ is an orthonormal basis for $\mathbb{R}^N$ and that Φ is a δ-stable embedding of $(\Psi(\Sigma_{2K_S}), \mathcal{R}(\Psi_I))$, where Ψ_I is an $N \times K_I$ submatrix of Ψ. Set $\Omega = \Phi\Psi_I$ and define P_Ω and $P_{\Omega^\perp}$ as in (10.2) and (10.3). Then $P_{\Omega^\perp}\Phi$ is a $\delta/(1-\delta)$-stable embedding of $(P_{\mathcal{R}(\Psi_I)^\perp}\Psi(\Sigma_{2K_S}), \{0\})$.*

Proof. This follows from the observation that $P_{\mathcal{R}(\Psi_I)^\perp}\Psi(\Sigma_{2K_S}) \subset \Psi(\Sigma_{2K_S})$ and then applying Corollary 10.1. □

We emphasize that in the above Corollary, $P_{\mathcal{R}(\Psi_I)^\perp}\Psi(\Sigma_{2K_S})$ will simply be the original family of sparse signals but with zeros in positions indexed by Ψ_I. One can easily verify that if $\delta \leq (\sqrt{2}-1)/\sqrt{2}$, then $\delta/(1-\delta) \leq \sqrt{2}-1$, and thus Corollary 10.2 is sufficient to ensure that the conditions for Theorem 3.2 are satisfied. We therefore conclude that under a slightly more restrictive bound on the required RIP constant, we can directly recover a sparse signal of interest x_S that is orthogonal to the interfering x_I without actually recovering x_I. Note that in addition to filtering out true interference, this framework is also relevant to the problem of signal recovery when the support is partially known, in which case the known support defines a subspace that can be thought of as interference to be rejected prior to recovering the remaining signal. Thus, our approach provides an alternative method for solving and analyzing the problem of CS recovery with partially known support considered in [144]. Furthermore, this result can also be useful in analyzing iterative recovery algorithms, as was demonstrated in Chapter 6, or in the case where we wish to recover a slowly varying signal as it evolves in time, as in [145].

This *cancel-then-recover* approach to signal recovery has a number of advantages. Observe that if we attempt to first recover x and then cancel x_I, then we require the RIP of order $2(K_S + K_I)$ to ensure that the *recover-then-cancel* approach will be successful. In contrast, filtering out x_I followed by recovery of x_S requires the RIP of order only $2K_S + K_I$. In certain cases (when K_I is significantly larger than K_S), this results in a substantial decrease in the required number of measurements. Furthermore, since all recovery algorithms have computational complexity that is at least linear in the sparsity of the recovered signal, this can also result in substantial computational savings for signal recovery.

10.1.3 Simulations and discussion

In this section we evaluate the performance of the cancel-then-recover approach suggested by Corollary 10.2. Rather than ℓ_1-minimization we use the CoSaMP algorithm since it more naturally lends itself towards a simple modification described below. More specifically, we evaluate three interference cancellation approaches.

1. **Cancel-then-recover:** This is the approach advocated in this section. We cancel out the contribution of x_I to the measurements y and directly recover x_S using the CoSaMP algorithm.

2. **Modified recovery:** Since we know the support of x_I, rather than cancelling out the contribution from x_I to the measurements, we modify a greedy algorithm such as CoSaMP to exploit the fact that part of the support of x is known in advance. This modification is made simply by forcing CoSaMP to always keep the elements of $\Lambda_I = \text{supp}(x_I)$ in the active set at each iteration. Essentially, this algorithm is exactly the same as the standard CoSaMP algorithm, but where we change the definition of $\text{hard}(x, K)$ to

$$[\text{hard}(x,K)]_i = \begin{cases} x_i, & |x_i| \text{ is among the } K \text{ largest elements of } |x| \text{ or } i \in \Lambda_I; \\ 0, & \text{otherwise.} \end{cases}$$

   After recovering $\widehat{x}$, we then set $\widehat{x}_n = 0$ for $n \in \Lambda_I$ to filter out the interference.

3. **Recover-then-cancel:** In this approach, we ignore the fact that we know the support of x_I and try to recover the signal x using the standard CoSaMP algorithm, and then set $\widehat{x}_n = 0$ for $n \in \Lambda_I$ as before.

In our experiments, we set $N = 1000$, $M = 200$, and $K_S = 10$. We then considered values of K_I from 1 to 100. We choose $\mathcal{S}_S$ and $\mathcal{S}_I$ by selecting random, non-overlapping sets of indices, so in this experiment, $\mathcal{S}_S$ and $\mathcal{S}_I$ are orthogonal (although they need

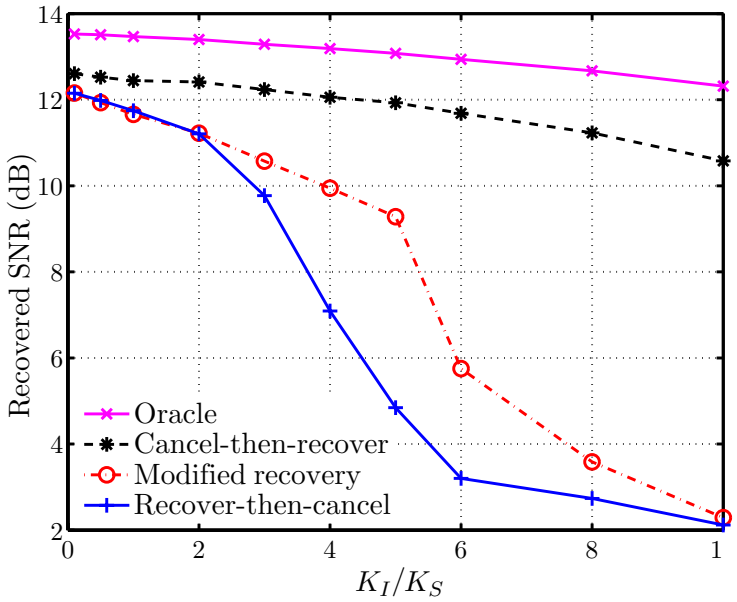

Figure 10.1: SNR of x_S recovered using the three different cancellation approaches for different ratios of K_I to K_S compared to the performance of an oracle.

not be in general, since $\widetilde{\mathcal{S}}_S$ will always be orthogonal to $\mathcal{S}_I$). For each value of K_I, we generated 2000 test signals where the coefficients were selected according to a Gaussian distribution and then contaminated with an N-dimensional Gaussian noise vector. For comparison, we also considered an oracle decoder that is given the support of both x_I and x_S and solves the least-squares problem restricted to the known support set.

We considered a range of signal-to-noise ratios (SNRs) and signal-to-interference ratios (SIRs). Figure 10.1 shows the results for the case where x_S and x_I are normalized to have equal energy (an SIR of 0dB) and where the variance of the noise is selected so that the SNR is 15dB. Our results were consistent for a wide range of SNR and SIR values, and we omit the plots due to space considerations.

Our results show that the cancel-then-recover approach performs significantly better than both of the other methods as K_I grows larger than K_S. In fact, the cancel-then-recover approach performs almost as well as the oracle decoder for the entire range of K_I. We also note that while the modified recovery method did perform

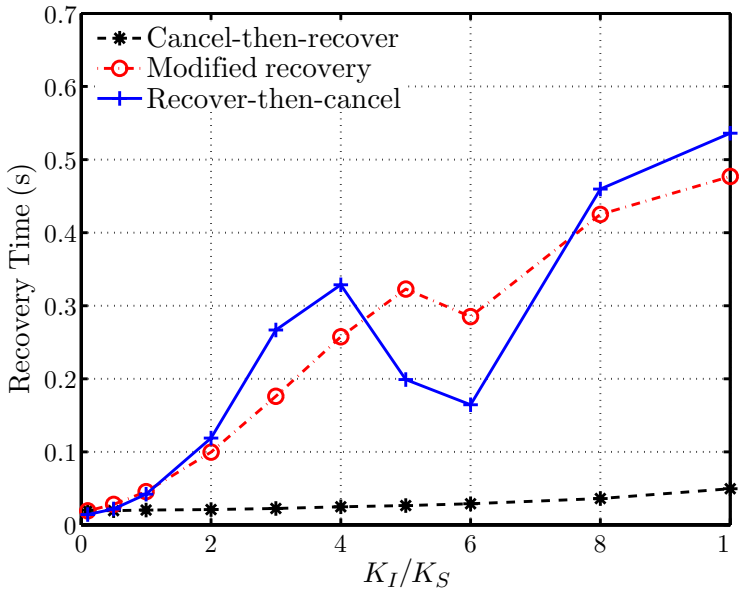

Figure 10.2: Recovery time for the three different cancellation approaches for different ratios of K_I to K_S.

slightly better than the recover-then-cancel approach, the improvement is relatively minor.

We observe similar results in Figure 10.2 for the recovery time (which includes the cost of computing P in the cancel-then-recover approach). The cancel-then-recover approach is significantly faster than the other approaches as K_I grows larger than K_S.

We also note that in the case where Φ admits a fast transform-based implementation (as is often the case for the constructions described in Chapter 5) the projections P_Ω and $P_{\Omega^\perp}$ can leverage the structure of Φ in order to ease the computational cost of applying P_Ω and $P_{\Omega^\perp}$. For example, Φ may consist of random rows of a Discrete Fourier Transform or a permuted Hadamard Transform matrix. In such a scenario, there are fast transform-based implementations of Φ and Φ^T. By observing that

$$P_\Omega = \Phi\Psi_I(\Psi_I^T\Phi^T\Phi\Psi_I)^{-1}\Psi_I^T\Phi^T$$

we see that one can use conjugate gradient methods to efficiently compute $P_\Omega y$ and, by extension, $P_{\Omega^\perp} y$ [44].

10.2 Bandstop Filtering

10.2.1 Filtering as subspace cancellation

The approach to "filtering" described above may seem somewhat foreign to someone more familiar with the classical notion of filtering. For example, there has been no mention thus far of frequency. We now show that these techniques can actually be applied to more classical filtering problems, and we specifically consider the problem of filtering out the contribution from a particular frequency band.

In order to do this, we model signals that live in the frequency band of interest as living in a subspace. To obtain a basis for this subspace, we will consider length-N windows of a bandlimited signal with band limits f_1 and f_2. Strictly speaking, such signals do *not* live in a subspace of $\mathbb{R}^N$, but one can show that they live very close to a low-dimensional subspace spanned by the first K *discrete prolate spheroidal sequences* (DPSS's) [146]. The DPSS's are the finite-length vectors that are simultaneously most concentrated in time and in frequency on the desired baseband bandwidth. In general, we can generate N DPSS's, but typically $K \ll N$ is sufficient to capture most of the energy in a bandlimited function. While there do exist rules of thumb for setting K, we will leave K as a parameter to be set by the user, with larger K allowing for better suppression of the undesired signal but also leading to slightly more distortion of the desired signal. We will let Ψ denote the $K \times N$ DPSS basis, which is generated by modulating baseband DPSS's by a cosine of frequency $(f_2 - f_1)/2$. If we have multiple bands we would like to filter out, then we can simply generate a basis for each and concatenate them into a matrix Ψ.

Once we have obtained the basis Ψ, interference cancellation can be performed as

described above by forming the matrix

$$P = I - \Phi\Psi(\Phi\Psi)^{\dagger}. \tag{10.10}$$

This will cancel signals supported on the bandwidth we wish to filter out while preserving the remainder of the signal.

10.2.2 Simulations and discussion

Our goal is to demonstrate the power of the proposed bandstop filtering algorithm. Towards this end, we first consider the problem of obtaining a rough estimate of the power spectral density (PSD) of a signal directly from the measurements. This will allow us to evaluate the effectiveness of our filtering approach. Let t denote the vector of time values corresponding to the sampling locations for the implicit Nyquist rate sampled version of x. We then consider a grid of possible frequency values $\{f_k\}_{k=1}^{k_{\max}}$ and compute

$$S(k) = |\langle \Phi e^{j2\pi f_k t}, y\rangle|^2 = |y^T \Phi e^{j2\pi f_k t}|^2 \tag{10.11}$$

for $k = 1, \ldots, k_{\max}$ where j here denotes $\sqrt{-1}$, i.e, we simply correlate y against a series of vectors that are the outputs of the random demodulator applied to pure sinusoid inputs. We could immediately provide a bound on the accuracy of this estimator using the results of Section 9.4, but we do not pursue this further.

Note that the filtering matrix P in (10.10) is an orthogonal projection. This is particularly useful, since after filtering we have $Py = P\Phi x$. One might expect that in order to use the proposed method for estimating the PSD, we would need to compute $P\Phi e^{j2\pi f_k t}$. However, since P is an orthogonal projection matrix we obtain

$$\langle Py, Pw\rangle = w^T P^T P y = w^T P^2 y = w^T P y = \langle Py, w\rangle.$$

Therefore, we can use the same PSD estimator *after* we have filtered the measurements without any modifications.

The result of our PSD estimation approach is a length-$k_{\max}$ vector S that provides an estimate of the power in the measurements at a sequence of frequencies of interest. This step is reminiscent of the first step of many greedy algorithms for CS recovery. Clearly, this sequence of frequencies should be sufficiently dense to be able to detect narrowband pulses. On the other hand, the spacing between frequencies should be relatively large in order to reduce the required number of computations. This allows for a tradeoff between the accuracy of the estimate S and the speed of its computation.

In the simulation results that follow, we consider the acquisition of a 300 MHz bandwidth in which 5 FM-modulated voice signals are transmitting at carrier frequencies unknown to the receiver. There is also a 140 kHz wide interferer at 200 MHz that is at least 35 dB stronger than any other signal. Each signal occupies roughly 12 kHz of bandwidth. The random demodulator compressively samples at a rate of 30 MHz (20 times undersampled).

Our experiments assume an input of a tuned and downconverted RF signal. Additionally, we assume the signal is collected in time-limited blocks, that the captured signals are known to be FM modulated, that there are a known number of signals in the collection bandwidth, and that the signals are separated by at least 30 kHz.

In Figure 10.3 we show the estimated PSDs before and after several stages of filtering. In *(a)* we show the true PSD of the original input signal, while in *(b)* we show the estimate obtained from compressive measurements using the method described above. Note that in *(b)* the effect of the large interferer dominates the PSD. However when we cancel it first (see *(c)*), we are able to obtain a much more accurate estimate of the smaller signal components. As shown, our original SNR of 40 dB is reduced to $\approx$ 25 dB, approximately a 15 dB loss. This is expected, since as described in Chapter 7 undersampling by a factor of 20 should result in an SNR

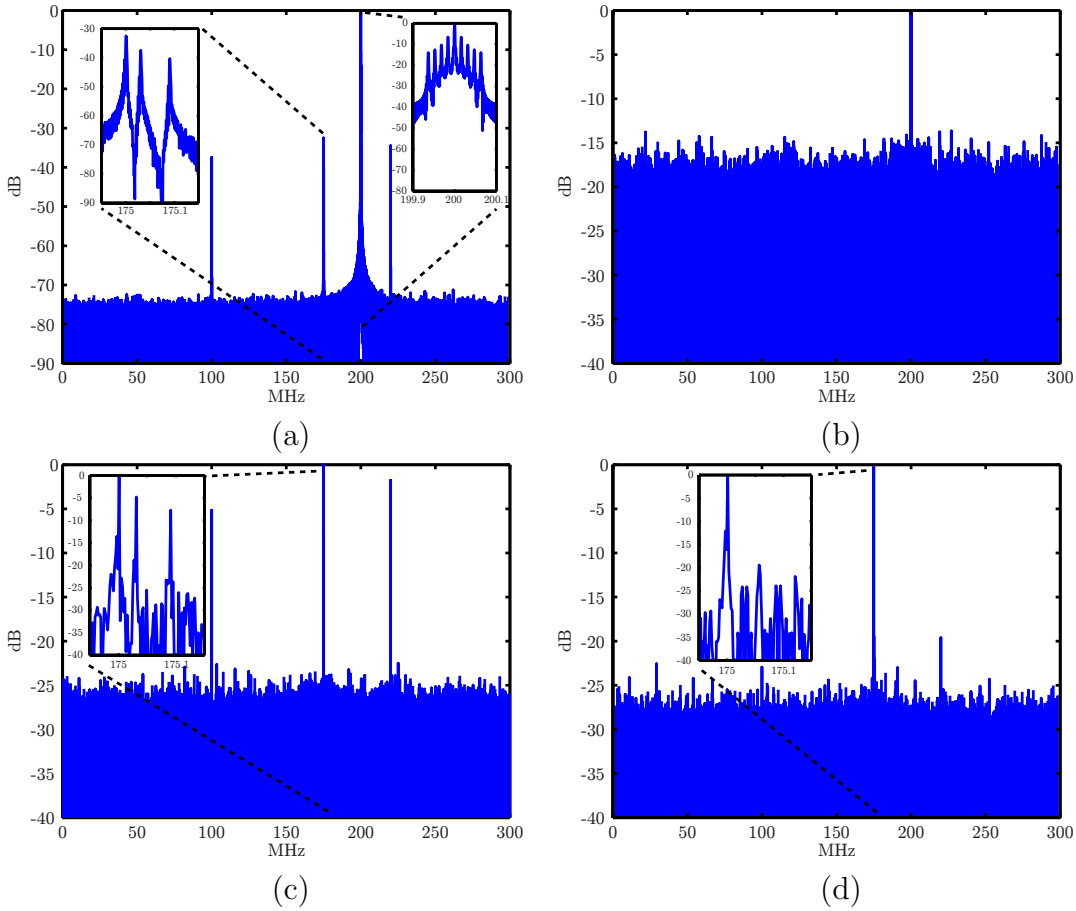

Figure 10.3: Normalized power spectral densities (PSDs). (a) PSD of original Nyquist-rate signal. (b) PSD estimated from compressive measurements. (c) PSD estimated from compressive measurements after filtering out interference. (d) PSD estimated from compressive measurements after further filtering.

loss of $3\log_2(20) \approx 13$ dB. In *(d)* the signal at 175 MHz is kept and the other signals are filtered out. In light of these results, in our view the bandstop filtering approach works remarkably well. While a number of further questions remaining concerning the limits of this approach to compressive filtering, this represents a very promising initial step.

Chapter 11

Conclusion

This thesis builds on the field of CS and illustrates how sparsity can be exploited to design efficient signal processing algorithms at all stages of the information processing pipeline, with a particular emphasis on the manner in which randomness can be exploited to design new kinds of acquisition systems for sparse signals. Specifically, our focus has been an analysis of the use of randomness in the design of compressive measurement systems that, when coupled with the appropriate signal models, enable robust and accurate recovery, as well as efficient algorithms for compressive-domain processing and inference. Our specific contributions have included: exploration and analysis of the appropriate properties for *sparse signal acquisition systems*; insight into the useful properties of *random measurement schemes*; analysis of *orthogonal greedy algorithms* for recovering sparse signals from random measurements; an exploration of the *impact of noise*, both structured and unstructured, in the context of random measurements; and algorithms that *directly solve inference problems* in the compressive domain without resorting to full-scale signal recovery, both reducing the cost of signal acquisition and reducing the complexity of the post-acquisition processing. To conclude, we reflect on some important remaining open problems and possible directions for future research.

11.1 Low-Dimensional Signal Models

In this thesis, we have primarily focused on sparse signal models as a way of capturing the fact that many high-dimensional signals actually have a limited number of degrees of freedom, and hence can be acquired using a small number of measurements compared to their dimensionality. However, there are a variety of other low-dimensional signal models that can also be exploited using similar techniques to those discussed in this thesis and which may be even more useful in some settings.

For example, a model closely related to sparsity is the set of low-rank matrices:

$$\mathcal{X} = \{X \in \mathbb{R}^{N_1 \times N_2} : \text{rank}(X) \leq R\}. \tag{11.1}$$

The set $\mathcal{X}$ consists of matrices X such that $X = \sum_{k=1}^{R} \sigma_k u_k v_k^*$ where $\sigma_1, \sigma_2, \ldots, \sigma_R \geq 0$ are the singular values, and $u_1, u_2, \ldots, u_R \in \mathbb{R}^{N_1}$, $v_1, v_2, \ldots, v_R \in \mathbb{R}^{N_2}$ are the singular vectors. Rather than constraining the number of elements used to construct the signal, we are constraining the number of nonzero singular values. One can easily observe by counting the number of degrees of freedom in the singular value decomposition that this set has dimension $R(N_1 + N_2 - R)$, which for small R is significantly less than the number of entries in the matrix — $N_1 N_2$. Low-rank matrices arise in a variety of practical settings. For example, low-rank (Hankel) matrices correspond to low-order linear, time-invariant systems [147]. In many data embedding problems, such as sensor geolocation, the matrix of pairwise distances will typically have rank 2 or 3 [148, 149]. Finally, low-rank matrices arise naturally in the context of collaborative filtering systems such as the now-famous Netflix recommendation system [150] and the related problem of *matrix completion*, where a low-rank matrix is recovered from a small sampling of its entries [151–153].

Parametric or manifold models form another important class of low-dimensional signal models. These models arise in cases where *(i)* a K-dimensional parameter θ

can be identified that carries the relevant information about a signal and *(ii)* the signal $f(\theta) \in \mathbb{R}^N$ changes as a continuous (typically nonlinear) function of these parameters. Typical examples include a one-dimensional (1-D) signal shifted by an unknown time delay (parameterized by the translation variable), a recording of a speech signal (parameterized by the underlying phonemes being spoken), and an image of a 3-D object at an unknown location captured from an unknown viewing angle (parameterized by the 3-D coordinates of the object and its roll, pitch, and yaw) [154–156]. In these and many other cases, the signal class forms a nonlinear K-dimensional manifold in $\mathbb{R}^N$, i.e.,

$$\mathcal{X} = \{f(\theta) : \theta \in \Theta\}, \tag{11.2}$$

where Θ is the K-dimensional parameter space. Manifold-based methods for image processing have attracted considerable attention, particularly in the machine learning community, and can be applied to diverse applications including data visualization, classification, estimation, detection, control, clustering, and learning [156–164]. Low-dimensional manifolds have also been proposed as approximate models for a number of nonparametric signal classes such as images of human faces and handwritten digits [165–167]. It is important to note that manifold models are closely related to both sparse models and low-rank matrices; the set of low-rank matrices forms a $R(N_1 + N_2 - R)$-dimensional Riemannian manifold [168], and while a union of subspaces fails to satisfy certain technical requirements of a topological manifold (due to its behavior at the origin), the manifold viewpoint can still be useful in understanding sparse data [169].

Beyond these models, prior work in the linear algebra community on matrix completion [151] provides a variety of alternative low-dimensional models that might prove useful. Additionally, one can consider combining these models to form new "hybrid"

models. For example, in [170, 171] a matrix is modeled as the sum of a low-rank component with a sparse component. One could consider similar hybrid models such as a sparse or low-rank component plus a manifold-modeled component, or a combination of different manifold-modeled components. In fact, the latter model has close connections to the finite rate of innovation framework [28], which although closely related in spirit has yet to be clearly unified with the CS framework.

In the study of such low-dimensional signal models, it is often difficult to separate the study of the theoretical properties of a particular model from the study of the algorithms that exploit this model. Theoretical results often have algorithmic consequences, and the goal of efficient algorithms can also prompt important theoretical questions. Hence, the development of a signal model must be concurrent with the study of algorithms. For all of these models, there are a variety of open theoretical and algorithmic questions regarding signal acquisition, recovery, and processing.

11.2 Signal Acquisition

We have argued in this thesis that for any signal model $\mathcal{X}$, a core theoretical question concerns the design of an acquisition system $\Phi : \mathcal{X} \to \mathbb{R}^M$ that preserves the structure of $\mathcal{X}$. In Part II we discussed this problem for the case where $\mathcal{X} = \Sigma_K$ and briefly touched on a few other models, but for most of the signal models described in Section 11.1 this remains an open problem. Specifically, for each of these models, we would like to be able to answer questions such as: Given a low-dimensional model $\mathcal{X}$, how can we design Φ so that the number of measurements M is as small as possible while retaining a stable embedding of $\mathcal{X}$? What are the relevant properties of $\mathcal{X}$ that determine how to choose Φ? And, how large M must be?

This thesis addresses the case where $\mathcal{X}$ represents K-sparse data in $\mathbb{R}^N$, showing that random constructions of Φ can achieve $M = O(K \log(N/K))$, which is essentially

optimal. However, it has also been shown that if the sparse data also exhibits some additional structure — for example, if the nonzero coefficients form a connected tree or obey some kind of graphical model — then it is possible to eliminate the $\log(N/K)$ factor [94]. A thorough examination of when this is possible and how substantial the gains can be has yet to be performed. Furthermore, we speculate that many of the same techniques will yield insight into problem of stably measuring low-rank matrices. Finally, while some results on stable embeddings of smooth submanifolds of $\mathbb{R}^N$ have been established [172, 173], these are not known to be optimal, and little is known for the cases where the manifold is not smooth or where the manifold is embedded in an infinite-dimensional space such as $L_2[0, 1]$.

Note also that in the typical compressive sensing framework, it is assumed that Φ is fixed in advance and in particular is *non-adaptive* to the structure in the signal. While this can result in a simpler overall system, it is also enlightening to consider what can be gained by using *adaptive* measurement schemes. In the context of sparse signal acquisition, recent results have indicated that it is possible to acquire a signal using substantially fewer measurements if one allows the measurements to be sequentially adapted based on the previous measurements [174–176]. Alternatively, given a fixed budget of measurements, an adaptive scheme can acquire a signal with significantly higher fidelity in the presence of noise. However, these benefits have only been studied in the context of sparsity — important questions remain for the low-rank matrix and manifold model settings.

In this thesis we have also described two particular practical applications where we can actually design hardware that directly acquires compressive measurements of continuous-time signals or images. There are numerous additional applications of sparse, low-rank, and manifold models in medical and scientific imaging, geophysical data analysis, and digital communications. As research into these and other low-dimensional models develops, it will certainly lead to additional opportunities

and applications. For example, the single-pixel camera described in Chapter 5 can also exploit manifold models for classification-driven imaging [56, 135, 136]. Matrix completion will also likely have applications in hyperspectral imaging and video acquisition. Further research into these models will lead to the development of a variety of new sensing platforms and devices.

11.3 Signal Recovery

A distinguishing feature of CS is that the measurements may preserve the information about the signal without providing such a simple method for recovering the original signal. The same phenomenon arises in the context of matrix recovery/completion. These are just two examples of an emerging *computational sensing* framework that attempts to leverage Moore's Law to improve the efficiency of sensing systems at a cost of increased computational requirements. For such measurement schemes to obtain wide applicability, it is important to develop provably *accurate* and *efficient* algorithms for recovery.

In this thesis, we have described a number of ways to solve the recovery problem in the context of sparse data, broadly separated into optimization-based methods and iterative methods. Similarly, the original formulations of the recovery algorithms for matrix completion, in which the goal is to recover a low-rank matrix from a small sample of its entries, also relied on optimization (in this case minimizing the *nuclear norm*, or the ℓ_1-norm of the vector of singular values of the matrix) to recover the matrix [152, 153, 177]. As in CS, subsequent work demonstrated that there also exist powerful greedy algorithms for solving the matrix recovery problem [178]. Given the potentially massive size of the matrices considered in many matrix completion applications, the need for efficient algorithms is even more acute than in CS, and improving these algorithms remains an important active area of research. Even less

is known about recovery in the case of more general manifold models. While it is possible in principle [179], there are no known polynomial-time algorithms for the general recovery problem.

11.4 Signal Processing and Inference

In this thesis we have taken a first step in developing a framework for CSP, where random measurements are processed directly to solve certain signal processing tasks. The algorithms discussed in Chapters 9 and 10 represent a foundation for the development of future CSP algorithms. There remain many open questions in this problem area. First, in this thesis, sparsity is leveraged primarily to argue that we can use random measurements to enable the design of agnostic and flexible hardware. It remains unknown if sparsity can be exploited more directly to aid in other compressive inference tasks. Some initial steps in this direction suggest that exploiting sparsity more directly may be beneficial. For example, in [61] sparsity is explicitly leveraged to perform classification with very few random measurements.

Manifold models have also been exploited in [135] to perform manifold-based image classification. These results are extended to a sensor network setting that exploits shared manifold structure in [60]. While most work in low-rank matrix models has focused on the recovery problem, there is also significant potential for exploiting these models for inference as well. In general, given a low-dimensional signal model $\mathcal{X}$, we would like to know when and how can we exploit the structure of $\mathcal{X}$ to extract information directly from Φx. Does this allow us to acquire fewer measurements than would be required for signal recovery? Can we do this using less computation than simply recovering the signal and then applying standard algorithms? The answers to these questions lie at the heart of the CSP paradigm and will serve to illuminate the usefulness of compressive measurements in DSP.

Bibliography

[1] E. Whittaker. On the functions which are represented by the expansions of the interpolation theory. *Proc. Royal Soc. Edinburgh, Sec. A*, 35:181–194, 1915.

[2] H. Nyquist. Certain topics in telegraph transmission theory. *Trans. AIEE*, 47:617–644, 1928.

[3] V. Kotelnikov. On the carrying capacity of the ether and wire in telecommunications. In *Izd. Red. Upr. Svyazi RKKA*, Moscow, Russia, 1933.

[4] C. Shannon. Communication in the presence of noise. *Proc. Institute of Radio Engineers*, 37(1):10–21, 1949.

[5] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.

[6] R. Walden. Analog-to-digital converter survey and analysis. *IEEE J. Selected Areas Comm.*, 17(4):539–550, 1999.

[7] D. Healy. Analog-to-information: Baa #05-35, 2005. Available online at `http://www.darpa.mil/mto/solicitations/baa05-35/s/index.html`.

[8] D. Donoho. Denoising by soft-thresholding. *IEEE Trans. Inform. Theory*, 41(3):613–627, 1995.

[9] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, San Diego, CA, 1999.

[10] E. Candès and T. Tao. The Dantzig Selector: Statistical estimation when p is much larger than n. *Ann. Stat.*, 35(6):2313–2351, 2007.

[11] B. Olshausen and D. Field. Emergence of simple-cell receptive field properties by learning a sparse representation. *Nature*, 381:607–609, 1996.

[12] R. DeVore. Nonlinear approximation. *Acta Numerica*, 7:51–150, 1998.

[13] E. Candès and J. Romberg. Quantitative robust uncertainty principles and optimally sparse decompositions. *Found. Comput. Math.*, 6(2):227–254, 2006.

[14] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.

[15] E. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, 2006.

[16] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, 2006.

[17] E. Candès. Compressive sampling. In *Proc. Int. Congress of Math.*, Madrid, Spain, Aug. 2006.

[18] D. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

[19] R. Baraniuk. Compressive sensing. *IEEE Signal Processing Mag.*, 24(4):118–120, 124, 2007.

[20] R. Prony. Essai expérimental et analytique sur les lois de la Dilatabilité des fluides élastiques et sur celles de la Force expansive de la vapeur de l'eau et de la vapeur de l'alkool, à différentes températures. *J. de l'École Polytechnique,* Floréal et Prairial III, 1(2):24–76, 1795. R. Prony is Gaspard Riche, baron de Prony.

[21] C. Carathéodory. Über den variabilitätsbereich der koeffizienten von potenzreihen, die gegebene werte nicht annehmen. *Math. Ann.*, 64:95–115, 1907.

[22] C. Carathéodory. Über den variabilitätsbereich der fourierschen konstanten von positiven harmonischen funktionen. *Rend. Circ. Mat. Palermo*, 32:193–217, 1911.

[23] P. Feng and Y. Bresler. Spectrum-blind minimum-rate sampling and reconstruction of multiband signals. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Atlanta, GA, May 1996.

[24] Y. Bresler and P. Feng. Spectrum-blind minimum-rate sampling and reconstruction of 2-D multiband signals. In *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Zurich, Switzerland, Sept. 1996.

[25] P. Feng. *Universal spectrum blind minimum rate sampling and reconstruction of multiband signals.* PhD thesis, University of Illinois at Urbana-Champaign, Mar. 1997.

[26] R. Venkataramani and Y. Bresler. Further results on spectrum blind sampling of 2-D signals. In *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Chicago, IL, Oct. 1998.

[27] Y. Bresler. Spectrum-blind sampling and compressive sensing for continuous-index signals. In *Proc. Work. Inform. Theory and Applications (ITA)*, San Diego, CA, Jan. 2008.

[28] M. Vetterli, P. Marziliano, and T. Blu. Sampling signals with finite rate of innovation. *IEEE Trans. Signal Processing*, 50(6):1417–1428, 2002.

[29] A. Beurling. Sur les intégrales de Fourier absolument convergentes et leur application à une transformation fonctionelle. In *Proc. Scandinavian Math. Congress*, Helsinki, Finland, 1938.

[30] T. Blumensath and M. Davies. Iterative hard thresholding for compressive sensing. *Appl. Comput. Harmon. Anal.*, 27(3):265–274, 2009.

[31] E. Candès, M. Wakin, and S. Boyd. Enhancing sparsity by weighted ℓ_1 minimization. *J. Fourier Anal. Appl.*, 14(5-6):877–905, 2008.

[32] R. Chartrand and V. Staneva. Restricted isometry properties and nonconvex compressive sensing. *Inverse Problems*, 24(3):035020, 2008.

[33] A. Cohen, W. Dahmen, and R. DeVore. Instance optimal decoding by thresholding in compressed sensing. In *Int. Conf. Harmonic Analysis and Partial Differential Equations*, Madrid, Spain, Jun. 2008.

[34] G. Cormode and S. Muthukrishnan. Combinatorial algorithms for compressed sensing. In *Proc. Coll. Struc. Inform. Comm. Complexity (SIROCCO)*, Chester, UK, Jul. 2006.

[35] W. Dai and O. Milenkovic. Subspace pursuit for compressive sensing signal reconstruction. *IEEE Trans. Inform. Theory*, 55(5):2230–2249, 2009.

[36] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Comm. Pure Appl. Math.*, 57(11):1413–1457, 2004.

[37] D. Donoho, I. Drori, Y. Tsaig, and J.-L. Stark. Sparse solution of underdetermined linear equations by Stagewise Orthogonal Matching Pursuit. Preprint, 2006.

[38] D. Donoho and Y. Tsaig. Fast solution of ℓ_1 norm minimization problems when the solution may be sparse. *IEEE Trans. Inform. Theory*, 54(11):4789–4812, 2008.

[39] M. Figueiredo, R. Nowak, and S. Wright. Gradient projections for sparse reconstruction: Application to compressed sensing and other inverse problems. *IEEE J. Select. Top. Signal Processing*, 1(4):586–597, 2007.

[40] E. Hale, W. Yin, and Y. Zhang. A fixed-point continuation method for ℓ_1-regularized minimization with applications to compressed sensing. Technical Report TR07-07, Rice Univ., CAAM Dept., 2007.

[41] S. Muthukrishnan. Some algorithmic problems and results in compressed sensing. In *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sept. 2006.

[42] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. *Found. Comput. Math.*, 9(3):317–334, 2009.

[43] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. *IEEE J. Select. Top. Signal Processing*, 4(2):310–316, 2010.

[44] D. Needell and J. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Appl. Comput. Harmon. Anal.*, 26(3):301–321, 2009.

[45] J. Tropp and A. Gilbert. Signal recovery from partial information via orthogonal matching pursuit. *IEEE Trans. Inform. Theory*, 53(12):4655–4666, 2007.

[46] J. Tropp and S. Wright. Computational methods for sparse solution of linear inverse problems. *Proc. IEEE*, 98(6):948–958, 2010.

[47] M. Lustig, D. Donoho, and J. Pauly. Rapid MR imaging with compressed sensing and randomly under-sampled 3DFT trajectories. In *Proc. Annual Meeting of ISMRM*, Seattle, WA, May 2006.

[48] M. Lustig, J. Lee, D. Donoho, and J. Pauly. Faster imaging with randomly perturbed, under-sampled spirals and ℓ_1 reconstruction. In *Proc. Annual Meeting of ISMRM*, Miami, FL, May 2005.

[49] M. Lustig, J. Santos, J. Lee, D. Donoho, and J. Pauly. Application of compressed sensing for rapid MR imaging. In *Proc. Work. Struc. Parc. Rep. Adap. Signaux (SPARS)*, Rennes, France, Nov. 2005.

[50] J. Trzasko and A. Manduca. Highly undersampled magnetic resonance image reconstruction via homotopic ℓ_0-minimization. *IEEE Trans. Med. Imaging*, 28(1):106–121, 2009.

[51] S. Vasanawala, M. Alley, R. Barth, B. Hargreaves, J. Pauly, and M. Lustig. Faster pediatric MRI via compressed sensing. In *Proc. Annual Meeting Soc. Pediatric Radiology (SPR)*, Carlsbad, CA, Apr. 2009.

[52] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk. Beyond Nyquist: Efficient sampling of sparse, bandlimited signals. *IEEE Trans. Inform. Theory*, 56(1):520–544, 2010.

[53] J. Romberg. Compressive sensing by random convolution. *SIAM J. Imag. Sci.*, 2(4):1098–1128, 2009.

[54] J. Tropp, M. Wakin, M. Duarte, D. Baron, and R. Baraniuk. Random filters for compressive sampling and reconstruction. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Toulouse, France, May 2006.

[55] J. Treichler, M. Davenport, and R. Baraniuk. Application of compressive sensing to the design of wideband signal acquisition receivers. In *Proc. U.S./Australia Joint Work. Defense Apps. of Signal Processing (DASP)*, Lihue, Hawaii, Sept. 2009.

[56] M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, and R. Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Mag.*, 25(2):83–91, 2008.

[57] R. Robucci, L. Chiu, J. Gray, J. Romberg, P. Hasler, and D. Anderson. Compressive sensing on a CMOS separable transform image sensor. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Las Vegas, NV, Apr. 2008.

[58] R. Marcia, Z. Harmany, and R. Willett. Compressive coded aperture imaging. In *Proc. IS&T/SPIE Symp. Elec. Imag.: Comp. Imag.*, San Jose, CA, Jan. 2009.

[59] D. Baron, M. Duarte, S. Sarvotham, M. Wakin, and R. Baraniuk. Distributed compressed sensing of jointly sparse signals. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2005.

[60] M. Davenport, C. Hegde, M. Duarte, and R. Baraniuk. Joint manifolds for data fusion. To appear in *IEEE Trans. Image Processing*, Oct. 2010.

[61] M. Duarte, M. Davenport, M. Wakin, and R. Baraniuk. Sparse signal detection from incoherent projections. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Toulouse, France, May 2006.

[62] S. Muthukrishnan. *Data Streams: Algorithms and Applications*, volume 1 of *Found. Trends in Theoretical Comput. Science*. Now Publishers, Boston, MA, 2005.

[63] A. Gilbert and P. Indyk. Sparse recovery using sparse matrices. *Proc. IEEE*, 98(6):937–947, 2010.

[64] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, New York, NY, 1999.

[65] S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comp.*, 20(1):33–61, 1998.

[66] B. Logan. *Properties of High-Pass Signals*. PhD thesis, Columbia Universuty, 1965.

[67] H. Taylor, S. Banks, and J. McCoy. Deconvolution with the ℓ_1 norm. *Geophysics*, 44(1):39–52, 1979.

[68] S. Levy and P. Fullagar. Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution. *Geophysics*, 46(9):1235–1243, 1981.

[69] C. Walker and T. Ulrych. Autoregressive recovery of the acoustic impedance. *Geophysics*, 48(10):1338–1350, 1983.

[70] R. Tibshirani. Regression shrinkage and selection via the lasso. *J. Royal Statist. Soc B*, 58(1):267–288, 1996.

[71] G. Davis, S. Mallat, and Z. Zhang. Adaptive time-frequency decompositions. *SPIE J. Opt. Engin.*, 33(7):2183–2191, 1994.

[72] M. Efroymson. *Mathematical Methods for Digital Computers*, chapter Multiple Regression Analysis. Wiley, New York, NY, 1960.

[73] R. Hocking. The analysis and selection of variables in linear regression. *Biometrics*, 32(1):1–49, 1976.

[74] S. Verdú. *Multiuser Detection*. Cambridge Univ. Press, Cambridge, England, 1998.

[75] Y. Pati, R. Rezaifar, and P. Krishnaprasad. Orthogonal Matching Pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 1993.

[76] J. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Inform. Theory*, 50(10):2231–2242, 2004.

[77] E. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51(12):4203–4215, 2005.

[78] E. Candès. The restricted isometry property and its implications for compressed sensing. *Comptes rendus de l'Académie des Sciences, Série I*, 346(9-10):589–592, 2008.

[79] P. Wojtaszczyk. Stability and instance optimality for Gaussian measurements in compressed sensing. *Found. Comput. Math.*, 10(1):1–13, 2010.

[80] K. Do Ba, P. Indyk, E. Price, and D. Woodruff. Lower bounds for sparse recovery. In *Proc. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, Austin, TX, Jan. 2010.

[81] A. Garnaev and E. Gluskin. The widths of Euclidean balls. *Dokl. An. SSSR*, 277:1048–1052, 1984.

[82] B. Kashin. The widths of certain finite dimensional sets and classes of smooth functions. *Izvestia*, 41:334–351, 1977.

[83] A. Pinkus. n-widths and optimal recovery. In *Proc. AMS Symposia Appl. Math.: Approx. Theory*, New Orleans, LA, Jan. 1986.

[84] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Const. Approx.*, 28(3):253–263, 2008.

[85] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k-term approximation. *J. Amer. Math. Soc.*, 22(1):211–231, 2009.

[86] V. Buldygin and Y. Kozachenko. *Metric Characterization of Random Variables and Random Processes*. American Mathematical Society, Providence, RI, 2000.

[87] M. Ledoux. *The Concentration of Measure Phenomenon*. American Mathematical Society, Providence, RI, 2001.

[88] S. Dasgupta and A. Gupta. An elementary proof of the Johnson-Lindenstrauss lemma. Technical Report TR-99-006, Univ. of Cal. Berkeley, Comput. Science Division, Mar. 1999.

[89] D. Achlioptas. Database-friendly random projections. In *Proc. Symp. Principles of Database Systems (PODS)*, Santa Barbara, CA, May 2001.

[90] R. DeVore, G. Petrova, and P. Wojtaszczyk. Instance-optimality in probability with an ℓ_1-minimization decoder. *Appl. Comput. Harmon. Anal.*, 27(3):275–288, 2009.

[91] K. Ball. *An Elementary Introduction to Modern Convex Geometry*, volume 31 of *Flavors of Geometry*. MSRI Publ., Cambridge Univ. Press, Cambridge, England, 1997.

[92] W. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Proc. Conf. Modern Anal. and Prob.*, New Haven, CT, Jun. 1982.

[93] T. Blumensath and M. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Trans. Inform. Theory*, 55(4):1872–1882, 2009.

[94] R. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Inform. Theory*, 56(4):1982–2001, 2010.

[95] R. Baraniuk and M. Wakin. Random projections of smooth manifolds. *Found. Comput. Math.*, 9(1):51–77, 2009.

[96] P. Indyk and A. Naor. Nearest-neighbor-preserving embeddings. *ACM Trans. Algorithms*, 3(3):31, 2007.

[97] P. Agarwal, S. Har-Peled, and H. Yu. Embeddings of surfaces, curves, and moving points in euclidean space. In *Proc. Symp. Comput. Geometry*, Gyeongju, Korea, Jun. 2007.

[98] S. Dasgupta and Y. Freund. Random projection trees and low dimensional manifolds. In *Proc. ACM Symp. Theory of Comput.*, Victoria, BC, May 2008.

[99] D. Takhar, J. Laska, M. Wakin, M. Duarte, D. Baron, S. Sarvotham, K. Kelly, and R. Baraniuk. A new compressive imaging camera architecture using optical-domain compression. In *Proc. IS&T/SPIE Symp. Elec. Imag.: Comp. Imag.*, San Jose, CA, Jan. 2006.

[100] D. Takhar, J. Laska, M. Wakin, M. Duarte, D. Baron, K. Kelly, and R. Baraniuk. A compressed sensing camera: New theory and an implementation using digital micromirrors. In *Proc. IS&T/SPIE Symp. Elec. Imag.: Comp. Imag.*, San Jose, CA, Jan. 2006.

[101] M. Wakin, J. Laska, M. Duarte, D. Baron, S. Sarvotham, D. Takhar, K. Kelly, and R. Baraniuk. Compressive imaging for video representation and coding. In *Proc. Picture Coding Symp.*, Beijing, China, Apr. 2006.

[102] M. Wakin, J. Laska, M. Duarte, D. Baron, S. Sarvotham, D. Takhar, K. Kelly, and R. Baraniuk. An architecture for compressive imaging. In *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Atlanta, GA, Oct. 2006.

[103] R. Coifman, F. Geshwind, and Y. Meyer. Noiselets. *Appl. Comput. Harmon. Anal.*, 10:27–44, 2001.

[104] E. Candès and J. Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23(3):969–985, 2007.

[105] A. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. Strauss. Near-optimal sparse Fourier representations via sampling. In *Proc. ACM Symp. Theory of Comput.*, Montreal, Canada, May 2002.

[106] A. Gilbert, S. Muthukrishnan, and M. Strauss. Improved time bounds for near-optimal sparse Fourier representations. In *Proc. SPIE Optics Photonics: Wavelets*, San Diego, CA, Aug. 2005.

[107] M. Mishali and Y. Eldar. Blind multi-band signal reconstruction: Compressed sensing for analog signals. *IEEE Trans. Signal Processing*, 57(3):993–1009, 2009.

[108] S. Kirolos, J. Laska, M. Wakin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. Baraniuk. Analog-to-information conversion via random demodulation. In *Proc. IEEE Dallas Circuits and Systems Work. (DCAS)*, Dallas, TX, Oct. 2006.

[109] J. Laska, S. Kirolos, M. Duarte, T. Ragheb, R. Baraniuk, and Y. Massoud. Theory and implementation of an analog-to-information convertor using random demodulation. In *Proc. IEEE Int. Symposium on Circuits and Systems (ISCAS)*, New Orleans, LA, May 2007.

[110] W. Bajwa, J. Haupt, G. Raz, S. Wright, and R. Nowak. Toeplitz-structured compressed sensing matrices. In *Proc. IEEE Work. Stat. Signal Processing*, Madison, WI, Aug. 2007.

[111] M. Davenport and M. Wakin. Analysis of Orthogonal Matching Pursuit using the restricted isometry property. To appear in *IEEE Trans. Inform. Theory*, Sept. 2010.

[112] T. Cai, G. Xu, and J. Zhang. On recovery of sparse signals via ℓ_1 minimization. *IEEE Trans. Inform. Theory*, 55(7):3388–3397, 2009.

[113] M. Davenport, P. Boufounos, and R. Baraniuk. Compressive domain interference cancellation. In *Proc. Work. Struc. Parc. Rep. Adap. Signaux (SPARS)*, Saint-Malo, France, Apr. 2009.

[114] T. Cai, L. Wang, and G. Xu. New bounds for restricted isometry constants. To appear in *IEEE Trans. Inform. Theory*, 2010.

[115] H. Rauhut. On the impossibility of uniform sparse reconstruction using greedy methods. *Sampl. Theory in Signal and Image Process.*, 7(2):197–215, 2008.

[116] P. Bechler and P. Wojtaszczyk. Error estimates for orthogonal matching pursuit and random dictionaries. Preprint, Aug. 2009.

[117] E.-T. Liu and V. Temlyakov. Orthogonal super greedy algorithm and application in compressed sensing. Technical Report 2010:01, Univ. of South Carolina, IMI, Jan. 2010.

[118] E. Livshitz. On efficiency of Orthogonal Matching Pursuit. Preprint, Apr. 2010.

[119] T. Zhang. Sparse recovery with Orthogonal Matching Pursuit under RIP. Preprint, May 2010.

[120] R. Vaughan, N. Scott, and R. White. The theory of bandpass sampling. *IEEE Trans. Signal Processing*, 39(9):1973–1984, 1991.

[121] J. Laska, M. Davenport, and R. Baraniuk. Exact signal recovery from corrupted measurements through the pursuit of justice. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2009.

[122] M. Davenport, J. Laska, P. Boufouons, and R. Baraniuk. A simple proof that random matrices are democratic. Technical Report TREE 0906, Rice Univ., ECE Dept., Nov. 2009.

[123] J. Laska, P. Boufounos, M. Davenport, and R. Baraniuk. Democracy in action: Quantization, saturation, and compressive sensing. Preprint, 2009.

[124] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Machine Intell.*, 31(2):210–227, 2009.

[125] J. Wright and Y. Ma. Dense error correction via ℓ_1 minimization. *IEEE Trans. Inform. Theory*, 56(7):3540–3560, 2010.

[126] R. Carrillo, K. Barner, and T. Aysal. Robust sampling and reconstruction methods for compressed sensing. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Taipei, Taiwan, Apr. 2009.

[127] A. Calderbank and I. Daubechies. The pros and cons of democracy. *IEEE Trans. Inform. Theory*, 48(6):1721–1725, 2002.

[128] S. Güntürk. *Harmonic analysis of two problems in signal compression*. PhD thesis, Princeton University, Sept. 2000.

[129] E. Candès. Integration of sensing and processing. In *Proc. IMA Annual Program Year Work.*, Minneapolis, MN, Dec. 2005.

[130] J. Romberg and M. Wakin. Compressed sensing: A tutorial. In *Proc. IEEE Work. Stat. Signal Processing*, Madison, WI, Aug. 2007.

[131] F. Beutler. Error-free recovery of signals from irregularly spaced samples. *SIAM Rev.*, 8:328–335, 1966.

[132] A. Aldroubi and K. Gröchenig. Nonuniform sampling and reconstruction in shift-invariant spaces. *SIAM Rev.*, 43(4):585–620, 2001.

[133] M. Davenport, M. Wakin, and R. Baraniuk. Detection and estimation with compressive measurements. Technical Report TREE-0610, Rice Univ., ECE Dept., Oct. 2006.

[134] M. Davenport, P. Boufounos, M. Wakin, and R. Baraniuk. Signal processing with compressive measurements. *IEEE J. Select. Top. Signal Processing*, 4(2):445–460, 2010.

[135] M. Davenport, M. Duarte, M. Wakin, J. Laska, D. Takhar, K. Kelly, and R. Baraniuk. The smashed filter for compressive classification and target recognition. In *Proc. IS&T/SPIE Symp. Elec. Imag.: Comp. Imag.*, San Jose, CA, Jan. 2007.

[136] M. Duarte, M. Davenport, M. Wakin, J. Laska, D. Takhar, K. Kelly, and R. Baraniuk. Multiscale random projections for compressive classification. In *Proc. IEEE Int. Conf. Image Processing (ICIP)*, San Antonio, TX, Sept. 2007.

[137] J. Haupt and R. Nowak. Compressive sampling for signal detection. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Honolulu, HI, Apr. 2007.

[138] J. Haupt, R. Castro, R. Nowak, G. Fudge, and A. Yeh. Compressive sampling for signal classification. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2006.

[139] S. Kay. *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*. Prentice Hall, Upper Saddle River, NJ, 1998.

[140] L. Scharf. *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Addison-Wesley, Reading, MA, 1991.

[141] N. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions, Volume 1*. Wiley-Interscience, New York, NY, 1994.

[142] N. Alon, P. Gibbons, Y. Matias, and M. Szegedy. Tracking join and self-join sizes in limited storage. In *Proc. Symp. Principles of Database Systems (PODS)*, Philadelphia, PA, May 1999.

[143] H. Rauhut, K. Schnass, and P. Vandergheynst. Compressed sensing and redundant dictionaries. *IEEE Trans. Inform. Theory*, 54(5):2210–2219, 2008.

[144] N. Vaswani and W. Lu. Modified-CS: Modifying compressive sensing for problems with partially known support. In *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Seoul, Korea, Jun. 2009.

[145] N. Vaswani. Analyzing least squares and Kalman filtered compressed sensing. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Taipei, Taiwan, Apr. 2009.

[146] D. Slepian. Prolate spheroidal wave functions, Fourier analysis, and uncertainty. V – The discrete case. *Bell Systems Tech. J.*, 57:1371–1430, 1978.

[147] J. Partington. *An Introduction to Hankel Operators*. Cambridge University Press, Cambridge, England, 1988.

[148] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[149] A. So and Y. Ye. Theory of semidefinite programming for sensor network localization. *Math. Programming, Series A and B*, 109(2):367–384, 2007.

[150] D. Goldberg, D. Nichols, B. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Comm. of the ACM*, 35(12):61–70, 1992.

[151] L. Hogben, editor. *Handbook of Linear Algebra*. Discrete Mathematics and its Applications. Chapman & Hall / CRC, Boca Raton, FL, 2007.

[152] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772, 2009.

[153] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum rank solutions of matrix equations via nuclear norm minimization. To appear in *SIAM Rev.*, 2010.

[154] H. Lu. *Geometric Theory of Images*. PhD thesis, University of California, San Diego, 1998.

[155] D. Donoho and C. Grimes. Image manifolds which are isometric to Euclidean space. *J. Math. Imag. and Vision*, 23(1):5–24, 2005.

[156] M. Wakin, D. Donoho, H. Choi, and R. Baraniuk. The multiscale structure of non-differentiable image manifolds. In *Proc. SPIE Optics Photonics: Wavelets*, San Diego, CA, Aug. 2005.

[157] J. Costa and A. Hero. Geodesic entropic graphs for dimension and entropy estimation in manifold learning. *IEEE Trans. Signal Processing*, 52(8):2210–2221, 2004.

[158] M. Belkin and P. Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, 56:209–239, 2004.

[159] P. Niyogi. Manifold regularization and semi-supervised learning: Some theoretical analyses. Technical Report TR-2008-01, Univ. of Chicago, Comput. Science Dept., 2008.

[160] J. Tenenbaum, V.de Silva, and J. Landford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 2000.

[161] D. Donoho and C. Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proc. Natl. Acad. Sci.*, 100(10):5591–5596, 2003.

[162] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, 15(6):1373–1396, 2003.

[163] R. Coifman and M. Maggioni. Diffusion wavelets. *Appl. Comput. Harmon. Anal.*, 21(1):53–94, 2006.

[164] K. Weinberger and L. Saul. Unsupervised learning of image manifolds by semidefinite programming. *Int. J. Computer Vision*, 70(1):77–90, 2006.

[165] M. Turk and A. Pentland. Eigenfaces for recognition. *J. Cognitive Neuroscience*, 3(1):71–86, 1991.

[166] G. Hinton, P. Dayan, and M. Revow. Modelling the manifolds of images of handwritten digits. *IEEE Trans. Neural Networks*, 8(1):65–74, 1997.

[167] D. Broomhead and M. Kirby. The Whitney Reduction Network: A method for computing autoassociative graphs. *Neural Comput.*, 13:2595–2616, 2001.

[168] B. Vandereycken and S. Vandewalle. Riemannian optimization approach for computing low-rank solutions of Lyapunov equations. In *Proc. SIAM Conf. on Optimization*, Boston, MA, May 2008.

[169] M. Davenport and R. Baraniuk. Sparse geodesic paths. In *Proc. AAAI Fall Symp. on Manifold Learning*, Arlington, VA, Nov. 2009.

[170] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky. Rank-sparsity incoherence for matrix decomposition. Preprint, 2009.

[171] E. Candès, X. Li, Y. Ma, and J. Wright. Robust Principal Component Analysis? Preprint, 2009.

[172] M. Wakin and R. Baraniuk. Random projections of signal manifolds. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, Toulouse, France, May 2006.

[173] K. Clarkson. Tighter bounds for random projections of manifolds. In *Proc. Symp. Comp. Geometry*, College Park, MD, Jun. 2008.

[174] J. Haupt, R. Castro, and R. Nowak. Adaptive discovery of sparse signals in noise. In *Proc. Asilomar Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Oct. 2008.

[175] J. Haupt, R. Castro, and R. Nowak. Adaptive sensing for sparse signal recovery. In *Proc. IEEE Digital Signal Processing Work.*, Marco Island, FL, Jan. 2009.

[176] J. Haupt, R. Castro, and R. Nowak. Distilled sensing: Selective sampling for sparse signal recovery. In *Proc. Int. Conf. Art. Intell. Stat. (AISTATS)*, Clearwater Beach, FL, Apr. 2009.

[177] E. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inform. Theory*, 56(5):2053–2080, 2010.

[178] R. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Trans. Inform. Theory*, 56(6):2980–2998, 2010.

[179] M. Wakin. Manifold-based signal recovery and parameter estimation from compressive measurements. Preprint, 2008.