

# Manifold-Based Signal Recovery and Parameter Estimation from Compressive Measurements

Michael B. Wakin\*

September 25, 2008

## Abstract

A field known as Compressive Sensing (CS) has recently emerged to help address the growing challenges of capturing and processing high-dimensional signals and data sets. CS exploits the surprising fact that the information contained in a sparse signal can be preserved in a small number of compressive (or random) linear measurements of that signal. Strong theoretical guarantees have been established on the accuracy to which sparse or near-sparse signals can be recovered from noisy compressive measurements. In this paper, we address similar questions in the context of a different modeling framework. Instead of sparse models, we focus on the broad class of manifold models, which can arise in both parametric and non-parametric signal families. Building upon recent results concerning the stable embeddings of manifolds within the measurement space, we establish both deterministic and probabilistic instance-optimal bounds in  $\ell_2$  for manifold-based signal recovery and parameter estimation from noisy compressive measurements. In line with analogous results for sparsity-based CS, we conclude that much stronger bounds are possible in the probabilistic setting. Our work supports the growing empirical evidence that manifold-based models can be used with high accuracy in compressive signal processing.

**Keywords.** Manifolds, dimensionality reduction, random projections, Compressive Sensing, sparsity, signal recovery, parameter estimation, Johnson-Lindenstrauss lemma.

**AMS Subject Classification.** 53A07, 57R40, 62H12, 68P30, 94A12, 94A29.

## 1 Introduction

### 1.1 Concise signal models

A significant byproduct of the Information Age has been an explosion in the sheer quantity of raw data demanded from sensing systems. From digital cameras to mobile devices, scientific computing to medical imaging, and remote surveillance to signals intelligence, the size (or dimension)  $N$  of a typical desired signal continues to increase. Naturally, the dimension  $N$  imposes a direct burden on the various stages of the data processing pipeline, from the data acquisition itself to the subsequent transmission, storage, and/or analysis; and despite rapid and continual improvements in computer processing power, other bottlenecks do remain, such as communication bandwidth over wireless channels, battery power in remote sensors and handheld devices, and the resolution/bandwidth of analog-to-digital converters.

---

\*Division of Engineering, Colorado School of Mines. Email: mwakin@mines.edu. This research was partially supported by NSF Grant DMS-0603606 and DARPA Grant HR0011-08-1-0078. The content of this article does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

Fortunately, in many cases, the information contained within a high-dimensional signal actually obeys some sort of concise, low-dimensional model. Such a signal may be described as having just  $K \ll N$  degrees of freedom for some  $K$ . Periodic signals bandlimited to a certain frequency are one example; they live along a fixed  $K$ -dimensional linear subspace of  $\mathbb{R}^N$ . Piecewise smooth signals are an example of *sparse signals*, which can be written as a succinct linear combination of just  $K$  elements from some basis such as a wavelet dictionary. Still other signals may live along  $K$ -dimensional submanifolds of the ambient signal space  $\mathbb{R}^N$ ; examples include collections of signals observed from multiple viewpoints in a camera or sensor network. In general, the conciseness of these models suggests the possibility for efficient processing and compression of these signals.

## 1.2 Compressive measurements

Recently, the conciseness of certain signal models has led to the use of *compressive measurements* for simplifying the data acquisition process. Rather than designing a sensor to measure a signal  $x \in \mathbb{R}^N$ , for example, it often suffices to design a sensor that can measure a much shorter vector  $y = \Phi x$ , where  $\Phi$  is a linear measurement operator represented as an  $M \times N$  matrix, and where typically  $M \ll N$ . As we discuss below in the context of Compressive Sensing (CS), when  $\Phi$  is properly designed, the requisite number of measurements  $M$  typically scales with the information level  $K$  of the signal, rather than with its ambient dimension  $N$ .

Surprisingly, the requirements on the measurement matrix  $\Phi$  can often be met by choosing  $\Phi$  randomly from an acceptable distribution. One distribution allows the entries of  $\Phi$  to be chosen as i.i.d. Gaussian random variables; another dictates that  $\Phi$  has orthogonal rows that span a random  $M$ -dimensional subspace of  $\mathbb{R}^N$ .

Physical architectures have been proposed for hardware that will enable the acquisition of signals using compressive measurements [11, 21, 24, 28]. The potential benefits for data acquisition are numerous. These systems can enable simple, low-cost acquisition of a signal directly in compressed form without requiring knowledge of the signal structure in advance. Some of the many possible applications include distributed source coding in sensor networks [4], medical imaging [29], high-rate analog-to-digital conversion [11, 24, 28], and error control coding [8].

## 1.3 Signal understanding from compressive measurements

Having acquired a signal  $x$  in compressed form (in the form of a measurement vector  $y$ ), there are many questions that may then be asked of the signal. These include:

- Q1. *Recovery*: What was the original signal  $x$ ?
- Q2. *Sketching*: Supposing that  $x$  was sparse or nearly so, what were the  $K$  basis vectors used to generate  $x$ ?
- Q3. *Parameter estimation*: Supposing  $x$  was generated from a  $K$ -dimensional parametric model, what was the original  $K$ -dimensional parameter that generated  $x$ ?

Given only the measurements  $y$  (possibly corrupted by noise), solving any of the above problems requires exploiting the concise,  $K$ -dimensional structure inherent in the signal.<sup>1</sup> CS addresses questions Q1 and Q2 under the assumption that the signal  $x$  is  $K$ -sparse (or approximately so) in some basis or dictionary; in Section 2 we outline several key theoretical bounds from CS regarding the accuracy to which these questions may be answered.

---

<sup>1</sup>Other problems, such as finding the nearest neighbor to  $x$  in a large database of signals [27], can also be solved using compressive measurements and do not require assumptions about the concise structure in  $x$ .

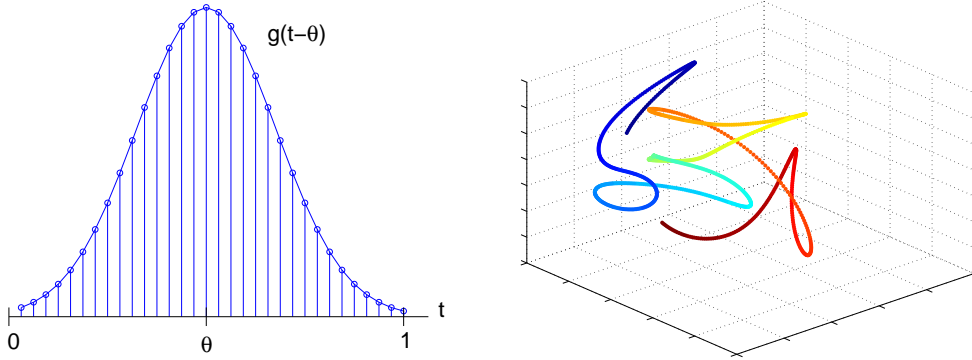


Figure 1: (a) The articulated signal  $f_\theta(t) = g(t - \theta)$  is defined via shifts of a primitive function  $g$ , where  $g$  is a Gaussian pulse. Each signal is sampled at  $N$  points, and as  $\theta$  changes, the resulting signals trace out a 1-D manifold in  $\mathbb{R}^N$ . (b) Projection of the manifold from  $\mathbb{R}^N$  onto a random 3-D subspace; the color/shading represents different values of  $\theta \in [0, 1]$ .

#### 1.4 Manifold models for signal understanding

In this paper, we will address these questions in the context of a different modeling framework for concise signal structure. Instead of sparse models, we focus on the broad class of *manifold models*, which arise both in settings where a  $K$ -dimensional parameter  $\theta$  controls the generation of the signal and also in non-parametric settings.

As a very simple illustration, consider the articulated signal in Figure 1(a). We let  $g(t)$  be a fixed continuous-time Gaussian pulse centered at  $t = 0$  and consider a shifted version of  $g$  denoted as the parametric signal  $f_\theta(t) := g(t - \theta)$  with  $t, \theta \in [0, 1]$ . We then suppose the discrete-time signal  $x = x_\theta \in \mathbb{R}^N$  arises by sampling the continuous-time signal  $f_\theta(t)$  uniformly in time, i.e.,  $x_\theta(n) = f_\theta(n/N)$  for  $n = 1, 2, \dots, N$ . As the parameter  $\theta$  changes, the signals  $x_\theta$  trace out a continuous one-dimensional (1-D) curve  $\mathcal{M} = \{x_\theta : \theta \in [0, 1]\} \subset \mathbb{R}^N$ . The conciseness of our model (in contrast with the potentially high dimension  $N$  of the signal space) is reflected in the low dimension of the path  $\mathcal{M}$ .

In the real world, manifold models may arise in a variety of settings. A  $K$ -dimensional parameter  $\theta$  could reflect uncertainty about the 1-D timing of the arrival of a signal (as in Figure 1(a)), the 2-D orientation and position of an edge in an image, the 2-D translation of an image under study, the multiple degrees of freedom in positioning a camera or sensor to measure a scene, the physical degrees of freedom in an articulated robotic or sensing system, or combinations of the above. Manifolds have also been proposed as approximate models for signal databases such as collections of images of human faces or of handwritten digits [5, 26, 35].

Consequently, the potential applications of manifold models are numerous in signal processing. In some applications, the signal  $x$  itself may be the object of interest, and the concise manifold model may facilitate the acquisition or compression of that signal. Alternatively, in parametric settings one may be interested in using a signal  $x = x_\theta$  to infer the parameter  $\theta$  that generated that signal. In an application known as manifold learning, one may be presented with a collection of data  $\{x_{\theta_1}, x_{\theta_2}, \dots, x_{\theta_n}\}$  sampled from a parametric manifold and wish to discover the underlying parameterization that generated that manifold. Multiple manifolds can also be considered simultaneously, for example in problems that require recognizing an object from one of  $n$  possible classes, where the viewpoint of the object is uncertain during the image capture process. In this case, we may wish to know which of  $n$  manifolds is closest to the observed image  $x$ .

While any of these questions may be answered with full knowledge of the high-dimensional signal

$x \in \mathbb{R}^N$ , there is growing theoretical and experimental support that they can also be answered from only compressive measurements  $y = \Phi x$ . In a recent paper, we have shown that given a sufficient number  $M$  of random measurements, one can ensure with high probability that a manifold  $\mathcal{M} \subset \mathbb{R}^N$  has a stable embedding in the measurement space  $\mathbb{R}^M$  under the operator  $\Phi$ , such that pairwise Euclidean and geodesic distances are approximately preserved on its image  $\Phi\mathcal{M}$ . We restate the precise result in Section 3, but a key aspect is that the number of requisite measurements  $M$  is linearly proportional to the information level of the signal, i.e., the dimension  $K$  of the manifold.

As a very simple illustration of this embedding phenomenon, Figure 1(b) presents an experiment where just  $M = 3$  compressive measurements are acquired from each point  $x_\theta$  described in Figure 1(a). We let  $N = 1024$  and construct a randomly generated  $3 \times N$  matrix  $\Phi$  with orthogonal rows. Each point  $x_\theta$  from the original manifold  $\mathcal{M} \subset \mathbb{R}^{1024}$  maps to a unique point  $\Phi x_\theta$  in  $\mathbb{R}^3$ ; the manifold embeds in the low-dimensional measurement space. Given any  $y = \Phi x_{\theta'}$  for  $\theta'$  unknown, then, it is possible to infer the value  $\theta'$  using only knowledge of the parametric model for  $\mathcal{M}$  and the measurement operator  $\Phi$ . Moreover, as the number  $M$  of compressive measurements increases, the manifold embedding becomes much more stable and remains highly self-avoiding.

Indeed, there is strong empirical evidence that, as a consequence of this phenomenon, questions such as Q1 (signal recovery) and Q3 (parameter estimation) can be accurately solved using only compressive measurements of a signal  $x$ , and that these procedures are robust to noise and to deviations of the signal  $x$  away from the manifold  $\mathcal{M}$  [15, 36]. Additional theoretical and empirical justification has followed for the manifold learning [25] and multiclass recognition problems [15] described above. Consequently, many of the advantages of compressive measurements that are beneficial in sparsity-based CS (low-cost sensor design, reduced transmission requirements, reduced storage requirements, lack of need for advance knowledge of signal structure, simplified computation in the low-dimensional space  $\mathbb{R}^M$ , etc.) may also be enjoyed in settings where manifold models capture the concise signal structure. Moreover, the use of a manifold model can often capture the structure of a signal in many fewer degrees of freedom  $K$  than would be required in any sparse representation, and thus the measurement rate  $M$  can be greatly reduced compared to sparsity-based CS approaches.

In this paper, we will focus on questions Q1 (signal recovery) and Q3 (parameter estimation) and reinforce the existing empirical work by establishing theoretical bounds on the accuracy to which these questions may be answered. We will consider both deterministic and probabilistic instance-optimal bounds, and we will see strong similarities to analogous results that have been derived for sparsity-based CS. As with sparsity-based CS, we show for manifold-based CS that for any fixed  $\Phi$ , uniform deterministic  $\ell_2$  recovery bounds for recovery of all  $x$  are necessarily poor. We then show that, as with sparsity-based CS, providing for any  $x$  a probabilistic bound that holds over most  $\Phi$  is possible with the desired accuracy. We consider both noise-free and noisy measurement settings and compare our bounds with sparsity-based CS.

## 1.5 Paper organization

We begin in Section 2 with a brief review of CS topics, to set notation and to outline several key results for later comparison. In Section 3 we discuss manifold models in more depth, restate our previous bound regarding stable embeddings of manifolds, and formalize our criteria for answering questions Q1 and Q3 in the context of manifold models. In Section 4, we confront the task of deriving deterministic instance-optimal bounds in  $\ell_2$ . In Section 5, we consider instead probabilistic instance-optimal bounds in  $\ell_2$ . We conclude in Section 6 with a final discussion.

## 2 Sparsity-Based Compressive Sensing

### 2.1 Sparse models

The concise modeling framework used in Compressive Sensing (CS) is *sparsity*. Consider a signal  $x \in \mathbb{R}^N$  and suppose the  $N \times N$  matrix  $\Psi = [\psi_1 \ \psi_2 \ \cdots \ \psi_N]$  forms an orthonormal basis for  $\mathbb{R}^N$ . We say  $x$  is  $K$ -sparse in the basis  $\Psi$  if for  $\alpha \in \mathbb{R}^N$  we can write

$$x = \Psi\alpha,$$

where  $\|\alpha\|_0 = K < N$ . (The  $\ell_0$ -norm notation counts the number of nonzeros of the entries of  $\alpha$ .) In a sparse representation, the actual information content of a signal is contained exclusively in the  $K < N$  positions and values of its nonzero coefficients.

For those signals that are approximately sparse, we may measure their proximity to sparse signals as follows. We define  $\alpha_K \in \mathbb{R}^N$  to be the vector containing only the largest  $K$  entries of  $\alpha$ , with the remaining entries set to zero. Similarly, we let  $x_K = \Psi\alpha_K$ . It is then common to measure the proximity to sparseness using either  $\|\alpha - \alpha_K\|_1$  or  $\|\alpha - \alpha_K\|_2$  (the latter of which equals  $\|x - x_K\|_2$  because  $\Psi$  is orthonormal).

### 2.2 Compressive measurements

CS uses the concept of sparsity to simplify the data acquisition process. Rather than designing a sensor to measure a signal  $x \in \mathbb{R}^N$ , for example, it often suffices to design a sensor that can measure a much shorter vector  $y = \Phi x$ , where  $\Phi$  is a linear measurement operator represented as an  $M \times N$  matrix, and typically  $M \ll N$ .

The measurement matrix  $\Phi$  must have certain properties in order to be suitable for CS. One desirable property (which leads to the theoretical results we mention in Section 2.3) is known as the Restricted Isometry Property (RIP) [7, 9, 10]. We say a matrix  $\Phi$  meets the *RIP of order  $K$  with respect to the basis  $\Psi$*  if for some  $\delta_K > 0$ ,

$$(1 - \delta_K) \|\alpha\|_2 \leq \|\Phi\Psi\alpha\|_2 \leq (1 + \delta_K) \|\alpha\|_2$$

holds for all  $\alpha \in \mathbb{R}^N$  with  $\|\alpha\|_0 \leq K$ . Intuitively, the RIP can be viewed as guaranteeing a *stable embedding* of the collection of  $K$ -sparse signals within the measurement space  $\mathbb{R}^M$ . In particular, supposing the RIP of order  $2K$  is satisfied with respect to the basis  $\Psi$ , then for all pairs of  $K$ -sparse signals  $x_1, x_2 \in \mathbb{R}^N$ , we have

$$(1 - \delta_{2K}) \|x_1 - x_2\|_2 \leq \|\Phi x_1 - \Phi x_2\|_2 \leq (1 + \delta_{2K}) \|x_1 - x_2\|_2. \quad (1)$$

Although deterministic constructions of matrices meeting the RIP are still a work in progress, it is known that the RIP often be met by choosing  $\Phi$  randomly from an acceptable distribution. For example, let  $\Psi$  be a fixed orthonormal basis for  $\mathbb{R}^N$  and suppose that

$$M \geq C_0 K \log(N/K) \quad (2)$$

for some constant  $C_0$ . Then supposing that the entries of the  $M \times N$  matrix  $\Phi$  are drawn as independent, identically distributed Gaussian random variables with mean 0 and variance  $\frac{1}{M}$ , it follows that with high probability  $\Phi$  meets the RIP of order  $K$  with respect to the basis  $\Psi$ . Two aspects of this construction deserve special notice: first, the number  $M$  of measurements required is linearly proportional to the information level  $K$ , and second, neither the sparse basis  $\Psi$  nor the locations of the nonzero entries of  $\alpha$  need be known when designing the measurement operator  $\Phi$ .

Other random distributions for  $\Phi$  may also be used, all requiring approximately the same number of measurements. One of these distributions [2, 14] dictates that  $\Phi = \sqrt{N/M} \Xi$ , where  $\Xi$  is an  $M \times N$  matrix having orthonormal rows that span a random  $M$ -dimensional subspace of  $\mathbb{R}^N$ . We refer to such choice of  $\Phi$  as a *random orthoprojector*.<sup>2</sup>

### 2.3 Signal recovery and sketching

Although the sparse structure of a signal  $x$  need not be known when collecting measurements  $y = \Phi x$ , a hallmark of CS is the use of the sparse model in order to facilitate understanding from the compressive measurements. A variety of algorithms have been proposed to answer Q1 (signal recovery), where we seek to solve the apparently undercomplete set of  $M$  linear equations  $y = \Phi x$  for  $N$  unknowns. The canonical method [6, 10, 17] is known as  $\ell_1$ -minimization and is formulated as follows: first solve

$$\hat{\alpha} = \arg \min_{\alpha' \in \mathbb{R}^N} \|\alpha'\|_1 \text{ subject to } y = \Phi \Psi \alpha', \quad (3)$$

and then set  $\hat{x} = \Psi \hat{\alpha}$ . Under this recovery program, the following bounds are known.

**Theorem 1** [12] *Suppose that  $\Phi$  satisfies the RIP of order  $2K$  with respect to  $\Psi$  and with constant  $\delta_{2K} < \sqrt{2} - 1$ . Let  $x \in \mathbb{R}^N$ , suppose  $y = \Phi x$ , and let the recovered estimates  $\hat{\alpha}$  and  $\hat{x}$  be as defined above. Then*

$$\|x - \hat{x}\|_2 = \|\alpha - \hat{\alpha}\|_2 \leq C_1 K^{-1/2} \|\alpha - \alpha_K\|_1 \quad (4)$$

for a constant  $C_1$ . In particular, if  $x$  is  $K$ -sparse, then  $\hat{x} = x$ .

This result can be extended to account for measurement noise.

**Theorem 2** [12] *Suppose that  $\Phi$  satisfies the RIP of order  $2K$  with respect to  $\Psi$  and with constant  $\delta_{2K} < \sqrt{2} - 1$ . Let  $x \in \mathbb{R}^N$ , and suppose that*

$$y = \Phi x + \eta$$

where  $\|\eta\|_2 \leq \epsilon$ . Then let

$$\hat{\alpha} = \arg \min_{\alpha' \in \mathbb{R}^N} \|\alpha'\|_1 \text{ subject to } \|y - \Phi \Psi \alpha'\|_2 \leq \epsilon,$$

and set  $\hat{x} = \Psi \hat{\alpha}$ . Then

$$\|x - \hat{x}\|_2 = \|\alpha - \hat{\alpha}\|_2 \leq C_1 K^{-1/2} \|\alpha - \alpha_K\|_1 + C_2 \epsilon. \quad (5)$$

for constants  $C_1$  (which is the same as above) and  $C_2$ .

These results are not unique to  $\ell_1$  minimization; similar bounds have been established for signal recovery using greedy iterative algorithms ROMP [32] and CoSAMP [31]. Bounds of this type are extremely encouraging for signal processing. From only  $M$  measurements, it is possible to recover  $x$  with quality that is comparable to its proximity to the nearest  $K$ -sparse signal, and if  $x$  itself is  $K$ -sparse and there is no measurement noise, then  $x$  can be recovered exactly. Moreover, despite the apparent ill-conditioning of the inverse problem, the measurement noise is not dramatically amplified in the recovery process.

---

<sup>2</sup>Our previous use of the term “random orthoprojector” in [3] excluded the normalization factor of  $\sqrt{N/M}$ . However we find it more appropriate to include this factor in the current paper.

These bounds are known as *deterministic, instance-optimal* bounds because they hold deterministically for any  $\Phi$  that meets the RIP, and because for a given  $\Phi$  they give a guarantee for recovery of any  $x \in \mathbb{R}^N$  based on its proximity to the concise model.

The use of  $\ell_1$  as a measure for proximity to the concise model (on the right hand side of (4) and (5)) arises due to the difficulty in establishing  $\ell_2$  bounds on the right hand side. Indeed, it is known that deterministic  $\ell_2$  instance-optimal bounds cannot exist that are comparable to (4) and (5). In particular, for any  $\Phi$ , to ensure that  $\|x - \hat{x}\|_2 \leq C_3 \|x - x_K\|_2$  for all  $x$ , it is known [13] that this requires that  $M \geq C_4 N$  regardless of  $K$ .

However, it is possible to obtain an instance-optimal  $\ell_2$  bound for sparse signal recovery in the noise-free setting by changing from a deterministic formulation to a *probabilistic* one [13, 16]. In particular, by considering any given  $x \in \mathbb{R}^N$ , it is possible to show that for *most* random  $\Phi$ , letting the measurements  $y = \Phi x$ , and recovering  $\hat{x}$  via  $\ell_1$ -minimization (3), it holds that

$$\|x - \hat{x}\|_2 \leq C_5 \|x - x_K\|_2. \quad (6)$$

While the proof of this statement [16] does not involve the RIP directly, it holds for many of the same random distributions that work for RIP matrices, and it requires the same number of measurements (2) up to a constant.

Similar bounds hold for the closely related problem of Q2 (sketching), where the goal is to use the compressive measurement vector  $y$  to identify and report only approximately  $K$  expansion coefficients that best describe the original signal, i.e., a sparse approximation to  $\alpha_K$ . In the case where  $\Psi = I$ , an efficient randomized measurement process coupled with a customized recovery algorithm [23] provides signal sketches that meet a deterministic mixed-norm  $\ell_2/\ell_1$  instance-optimal bound analogous to (4). A desirable aspect of this construction is that the computational complexity scales with only  $\log(N)$  (and is polynomial in  $K$ ); this is possible because only approximately  $K$  pieces of information must be computed to describe the signal. For signals that are sparse in the Fourier domain ( $\Psi$  consists of the DFT vectors), probabilistic  $\ell_2/\ell_2$  instance-optimal bounds have also been established [22] that are analogous to (6).

## 3 Compressive Measurements of Manifold-Modeled Signals

### 3.1 Manifold models

As we have discussed in Section 1.4, there are many possible modeling frameworks for capturing concise signal structure. Among these possibilities are the broad class of manifold models.

Manifold models arise, for example, in settings where the signals of interest vary continuously as a function of some  $K$ -dimensional parameter. Suppose, for instance, that there exists some parameter  $\theta$  that controls the generation of the signal. We let  $x_\theta \in \mathbb{R}^N$  denote the signal corresponding to the parameter  $\theta$ , and we let  $\Theta$  denote the  $K$ -dimensional parameter space from which  $\theta$  is drawn. In general,  $\Theta$  itself may be a  $K$ -dimensional manifold and need not be embedded in an ambient Euclidean space. For example, supposing  $\theta$  describes the 1-D rotation parameter in a top-down satellite image, we have  $\Theta = S^1$ .

Under certain conditions on the parameterization  $\theta \mapsto x_\theta$ , it follows that

$$\mathcal{M} := \{x_\theta : \theta \in \Theta\}$$

forms a  $K$ -dimensional *submanifold* of  $\mathbb{R}^N$ . An appropriate visualization is that the set  $\mathcal{M}$  forms a nonlinear  $K$ -dimensional “surface” within the high-dimensional ambient signal space  $\mathbb{R}^N$ . Depending on the circumstances, we may measure the distance between two points  $x_{\theta_1}$  and  $x_{\theta_2}$  on

the manifold  $\mathcal{M}$  using either the ambient Euclidean distance

$$\|x_{\theta_1} - x_{\theta_2}\|_2$$

or the geodesic distance along the manifold, which we denote as  $d_{\mathcal{M}}(x_{\theta_1}, x_{\theta_2})$ . In the case where the geodesic distance along  $\mathcal{M}$  equals the native distance in parameter space, i.e., when

$$d_{\mathcal{M}}(x_{\theta_1}, x_{\theta_2}) = d_{\Theta}(\theta_1, \theta_2), \tag{7}$$

we say that  $\mathcal{M}$  is *isometric* to  $\Theta$ . The definition of the distance  $d_{\Theta}(\theta_1, \theta_2)$  depends on the appropriate metric for the parameter space  $\Theta$ ; supposing  $\Theta$  is a convex subset of Euclidean space, then we have  $d_{\Theta}(\theta_1, \theta_2) = \|\theta_1 - \theta_2\|_2$ .

While our discussion above concentrates on the case of manifolds  $\mathcal{M}$  generated by underlying parameterizations, we stress that manifolds have also been proposed as approximate low-dimensional models within  $\mathbb{R}^N$  for nonparametric signal classes such as images of human faces or handwritten digits [5, 26, 35]. These signal families may also be considered.

The results we present in this paper will make reference to certain characteristic properties of the manifold under study. These terms are originally defined in [3, 33] and are repeated here for completeness. First, our results will depend on a measure of regularity for the manifold. For this purpose, we adopt the *condition number* defined recently by Niyogi et al. [33].

**Definition 1** [33] *Let  $\mathcal{M}$  be a compact Riemannian submanifold of  $\mathbb{R}^N$ . The condition number is defined as  $1/\tau$ , where  $\tau$  is the largest number having the following property: The open normal bundle about  $\mathcal{M}$  of radius  $r$  is embedded in  $\mathbb{R}^N$  for all  $r < \tau$ .*

The condition number  $1/\tau$  controls both local properties and global properties of the manifold. Its role is summarized in two key relationships [33]. First, the curvature of any unit-speed geodesic path on  $\mathcal{M}$  is bounded by  $1/\tau$ . Second, at long geodesic distances, the condition number controls how close the manifold may curve back upon itself. For example, supposing  $x_1, x_2 \in \mathcal{M}$  with  $d_{\mathcal{M}}(x_1, x_2) > \tau$ , it must hold that  $\|x_1 - x_2\|_2 > \tau/2$ .

We also require a notion of “geodesic covering regularity” for a manifold. While this property is not the focus of the present paper, we include its definition in Appendix A for completeness.

We conclude with a brief but concrete example to illustrate specific values for these quantities. Let  $N > 0$ ,  $\kappa > 0$ ,  $\Theta = \mathbb{R} \bmod 2\pi$ , and suppose  $x_{\theta} \in \mathbb{R}^N$  is given by

$$x_{\theta} = [\kappa \cos(\theta); \kappa \sin(\theta); 0; 0; \dots 0]^T.$$

In this case,  $\mathcal{M} = \{x_{\theta} : \theta \in \Theta\}$  forms a circle of radius  $\kappa$  in the  $x(1), x(2)$  plane. The manifold dimension  $K = 1$ , the condition number  $\tau = \kappa$ , and the geodesic covering regularity  $R$  can be chosen as any number larger than  $\frac{1}{2}$ . We also refer in our results to the  $K$ -dimensional volume  $V$  of the  $\mathcal{M}$ , which in this example corresponds to the circumference  $2\pi\kappa$  of the circle.

### 3.2 Stable embeddings of manifolds

In cases where the signal class of interest  $\mathcal{M}$  forms a low-dimensional submanifold of  $\mathbb{R}^N$ , we have theoretical justification that the information necessary to distinguish and recover signals  $x \in \mathcal{M}$  can be well-preserved under a sufficient number of compressive measurements  $y = \Phi x$ . In particular, we have recently shown that an RIP-like property holds for families of manifold-modeled signals.



**Theorem 3** [3] *Let  $\mathcal{M}$  be a compact  $K$ -dimensional Riemannian submanifold of  $\mathbb{R}^N$  having condition number  $1/\tau$ , volume  $V$ , and geodesic covering regularity  $R$ . Fix  $0 < \epsilon < 1$  and  $0 < \rho < 1$ . Let  $\Phi$  be a random  $M \times N$  orthoprojector with*

$$M = O\left(\frac{K \log(NVR\tau^{-1}\epsilon^{-1}) \log(1/\rho)}{\epsilon^2}\right). \quad (8)$$

*If  $M \leq N$ , then with probability at least  $1 - \rho$  the following statement holds: For every pair of points  $x_1, x_2 \in \mathcal{M}$ ,*

$$(1 - \epsilon) \|x_1 - x_2\|_2 \leq \|\Phi x_1 - \Phi x_2\|_2 \leq (1 + \epsilon) \|x_1 - x_2\|_2. \quad (9)$$

The proof of this theorem involves the Johnson-Lindenstrauss Lemma [1, 2, 14], which guarantees a stable embedding for a finite point cloud under a sufficient number of random projections. In essence, manifolds with higher volume or with greater curvature have more complexity and require a more dense covering for application of the Johnson-Lindenstrauss Lemma; this leads to an increased number of measurements (8).

By comparing (1) with (9), we see a strong analogy to the RIP of order  $2K$ . This theorem establishes that, like the class of  $K$ -sparse signals, a collection of signals described by a  $K$ -dimensional manifold  $\mathcal{M} \subset \mathbb{R}^N$  can have a stable embedding in an  $M$ -dimensional measurement space. Moreover, the requisite number of random measurements  $M$  is once again linearly proportional to the information level (or number of degrees of freedom)  $K$ .

As was the case with the RIP for sparse signal processing, this result has a number of possible implications for manifold-based signal processing. Individual signals obeying a manifold model can be acquired and stored efficiently using compressive measurements, and it is unnecessary to employ the manifold model itself as part of the compression process. Rather, the model need be used only for signal understanding from the compressive measurements. Problems such as Q1 (signal recovery) and Q3 (parameter estimation) can be addressed. We have reported promising experimental results with various classes of parametric signals [15, 36]. We have also extended Theorem 3 to the case of multiple manifolds that are simultaneously embedded [15]; this allows both the classification of an observed object to one of several possible models (different manifolds) and the estimation of a parameter within that class (position on a manifold). Moreover, collections of signals obeying a manifold model (such as multiple images of a scene photographed from different perspectives) can be acquired using compressive measurements, and the resulting manifold structure will be preserved among the suite of measurement vectors in  $\mathbb{R}^M$ . We have provided empirical and theoretical support for the use of manifold learning in the reduced-dimensional space [25]; this can dramatically simplify the computational and storage demands on a system for processing large databases of signals.

### 3.3 Signal recovery and parameter estimation

In this paper, we provide theoretical justification for the encouraging experimental results that have been observed for problems Q1 (signal recovery) and Q3 (parameter estimation).

To be specific, let us consider a length- $N$  signal  $x$  that, rather than being  $K$ -sparse, we assume lives on or near some known  $K$ -dimensional manifold  $\mathcal{M} \subset \mathbb{R}^N$ . From a collection of measurements

$$y = \Phi x + \eta,$$

where  $\Phi$  is a random  $M \times N$  matrix and  $\eta \in \mathbb{R}^M$  is an additive noise vector, we would like to recover either  $x$  or a parameter  $\theta$  that generates  $x$ .

For the signal recovery problem, we will consider the following as a method for estimating  $x$ :

$$\hat{x} = \arg \min_{x' \in \mathcal{M}} \|y - \Phi x'\|_2, \quad (10)$$

supposing here and elsewhere that the minimum is uniquely defined. We also let  $x^*$  be the optimal “nearest neighbor” to  $x$  on  $\mathcal{M}$ , i.e.,

$$x^* = \arg \min_{x' \in \mathcal{M}} \|x - x'\|_2. \quad (11)$$

To consider signal recovery successful, we would like to guarantee that  $\|x - \hat{x}\|_2$  is not much larger than  $\|x - x^*\|_2$ .

For the parameter estimation problem, where we presume  $x \approx x_\theta$  for some  $\theta \in \Theta$ , we propose a similar method for estimating  $\theta$  from the compressive measurements:

$$\hat{\theta} = \arg \min_{\theta' \in \Theta} \|y - \Phi x_{\theta'}\|_2. \quad (12)$$

Let  $\theta^*$  be the “optimal estimate” that could be obtained using the full data  $x \in \mathbb{R}^N$ , i.e.,

$$\theta^* = \arg \min_{\theta' \in \Theta} \|x - x_{\theta'}\|_2. \quad (13)$$

(If  $x = x_\theta$  exactly for some  $\theta$ , then  $\theta^* = \theta$ ; otherwise this formulation allows us to consider signals  $x$  that are not precisely on the manifold  $\mathcal{M}$  in  $\mathbb{R}^N$ . This generalization has practical relevance; a local image block, for example, may only approximately resemble a straight edge, which has a simple parameterization.) To consider parameter estimation successful, we would like to guarantee that  $d_\Theta(\hat{\theta}, \theta^*)$  is small.

As we will see, bounds pertaining to accurate signal recovery can often be extended to imply accurate parameter estimation as well. However, the relationships between distance  $d_\Theta$  in parameter space and distances  $d_{\mathcal{M}}$  and  $\|\cdot\|_2$  in the signal space can vary depending on the parametric signal model under study. Thus, for the parameter estimation problem, our ability to provide generic bounds on  $d_\Theta(\hat{\theta}, \theta^*)$  will be restricted. In this paper we focus primarily on the signal recovery problem and provide preliminary results for the parameter estimation problem that pertain most strongly to the case of isometric parameterizations.

In this paper, we do not confront in depth the question of how a recovery program such as (10) can be efficiently solved. Some discussion of this matter is provided in [3], with application-specific examples provided in [15, 36]. Unfortunately, it is difficult to propose a single general-purpose algorithm for solving (10) in  $\mathbb{R}^M$ , as even the problem (11) in  $\mathbb{R}^N$  may be difficult to solve depending on certain nuances (such as topology) of the individual manifold. Nonetheless, iterative algorithms such as Newton’s method [37] have proved helpful in many problems to date. Additional complications arise when the manifold  $\mathcal{M}$  is non-differentiable, as may happen when the signals  $x$  represent 2-D images. However, just as a multiscale regularization can be incorporated into Newton’s method for solving (11) (see [37]), an analogous regularization can be incorporated into a compressive measurement operator  $\Phi$  to facilitate Newton’s method for solving (10) (see [18, 36]). For manifolds that lack differentiability, additional care must be taken when applying results such as Theorem 3; we defer a study of these matters to a subsequent paper.

In the following, we will consider both deterministic and probabilistic instance-optimal bounds for signal recovery and parameter estimation, and we will draw comparisons to the sparsity-based

CS results of Section 2.3. Our bounds are formulated in terms of generic properties of the manifold (as mentioned in Section 3.1), which will vary from signal model to signal model. In some cases, calculating these may be possible, whereas in other cases it may not. Nonetheless, we feel the results in this paper highlight the relative importance of these properties in determining the requisite number of measurements. Finally, to simplify analysis we will focus on random orthoprojectors for the measurement operator  $\Phi$ , although our results may be extended to other random distributions such as the Gaussian [3].

## 4 A deterministic instance-optimal bound in $\ell_2$

We begin by seeking a deterministic instance-optimal bound. That is, for a measurement matrix  $\Phi$  that meets (9) for all  $x_1, x_2 \in \mathcal{M}$ , we seek an upper bound for the relative reconstruction error

$$\frac{\|x - \hat{x}\|_2}{\|x - x^*\|_2}$$

that holds uniformly for all  $x \in \mathbb{R}^N$ . In this section we consider only the signal recovery problem; however, similar bounds would apply to parameter estimation. We have the following result for the noise-free case, which applies not only to the manifolds described in Theorem 3 but also to more general sets.

**Theorem 4** *Let  $\mathcal{M} \subset \mathbb{R}^N$  be any subset of  $\mathbb{R}^N$ , and let  $\Phi$  denote an  $M \times N$  orthoprojector satisfying (9) for all  $x_1, x_2 \in \mathcal{M}$ . Suppose  $x \in \mathbb{R}^N$ , let  $y = \Phi x$ , and let the recovered estimate  $\hat{x}$  and the optimal estimate  $x^*$  be as defined in (10) and (11). Then*

$$\frac{\|x - \hat{x}\|_2}{\|x - x^*\|_2} \leq \sqrt{\frac{4N}{M(1-\epsilon)^2} - 3} + 2\sqrt{\frac{N}{M(1-\epsilon)^2} - 1}. \quad (14)$$

**Proof:** See Appendix B.

As  $\frac{M}{N} \rightarrow 0$ , the bound on the right hand side of (14) grows as  $\frac{2}{1-\epsilon}\sqrt{\frac{N}{M}}$ . Unfortunately, this is not desirable for signal recovery. Supposing, for example, that we wish to ensure  $\|x - \hat{x}\|_2 \leq C_6 \|x - x^*\|_2$  for all  $x \in \mathbb{R}^N$ , then using the bound (14) we would require that  $M \geq C_7 N$  regardless of the dimension  $K$  of the manifold.

The weakness of this bound is a geometric necessity; indeed, the bound itself is quite tight in general, as the following simple example illustrates. Suppose  $N \geq 2$  and let  $\mathcal{M}$  denote the line segment in  $\mathbb{R}^N$  joining the points  $(0, 0, \dots, 0)$  and  $(1, 0, 0, \dots, 0)$ . Let  $0 \leq \gamma < \pi/2$  for some  $\gamma$ , let  $M = 1$ , and let the  $1 \times N$  measurement matrix

$$\Phi = \sqrt{N} [\cos(\gamma); -\sin(\gamma); 0; 0; \dots; 0].$$

Any  $x_1 \in \mathcal{M}$  we may write as  $x_1 = (x_1(1), 0, 0, \dots, 0)$ , and it follows that  $\Phi x_1 = \sqrt{N} \cos(\gamma) x_1(1)$ . Thus for any pair  $x_1, x_2 \in \mathcal{M}$ , we have

$$\frac{\|\Phi x_1 - \Phi x_2\|_2}{\|x_1 - x_2\|_2} = \frac{|\sqrt{N} \cos(\gamma) x_1(1) - \sqrt{N} \cos(\gamma) x_2(1)|}{|x_1(1) - x_2(1)|} = \sqrt{N} \cos(\gamma).$$

We suppose that  $\cos(\gamma) < \frac{1}{\sqrt{N}}$  and thus referring to equation (9) we have  $(1-\epsilon) = \sqrt{N} \cos(\gamma)$ . Now, we may consider the signal  $x = (1, \tan(\pi/2 - \gamma), 0, 0, \dots, 0)$ . We then have that  $x^* = (1, 0, 0, \dots, 0)$ , and  $\|x - x^*\|_2 = \tan(\pi/2 - \gamma)$ . We also have that  $\Phi x = \sqrt{N}(\cos(\gamma) - \sin(\gamma) \tan(\pi/2 - \gamma)) = 0$ . Thus  $\hat{x} = (0, 0, \dots, 0)$  and  $\|x - \hat{x}\|_2 = \frac{1}{\cos(\pi/2 - \gamma)}$ , and so

$$\frac{\|x - \hat{x}\|_2}{\|x - x^*\|_2} = \frac{1}{\cos(\pi/2 - \gamma) \tan(\pi/2 - \gamma)} = \frac{1}{\sin(\pi/2 - \gamma)} = \frac{1}{\cos(\gamma)} = \frac{\sqrt{N}}{1 - \epsilon}.$$

It is worth recalling that, as we discussed in Section 2.3, similar difficulties arise in sparsity-based CS when attempting to establish a deterministic  $\ell_2$  instance-optimal bound. In particular, to ensure that  $\|x - \hat{x}\|_2 \leq C_3 \|x - x_K\|_2$  for all  $x \in \mathbb{R}^N$ , it is known [13] that this requires  $M \geq C_4 N$  regardless of the sparsity level  $K$ .

In sparsity-based CS, there have been at least two types of alternative approaches. The first are the deterministic “mixed-norm” results of the type given in (4) and (5). These involve the use of an alternative norm such as the  $\ell_1$  norm to measure the distance from the coefficient vector  $\alpha$  to its best  $K$ -term approximation  $\alpha_K$ . While it may be possible to pursue similar directions for manifold-modeled signals, we feel this is undesirable as a general approach because when sparsity is no longer part of the modeling framework, the  $\ell_1$  norm has less of a natural meaning. Instead, we prefer to seek bounds using  $\ell_2$ , as that is the most conventional norm used in signal processing to measure energy and error.

Thus, the second type of alternative bounds in sparsity-based CS have involved  $\ell_2$  bounds in probability, as we discussed in Section 2.3. Indeed, the performance of both sparsity-based and manifold-based CS is often much better in practice than a deterministic  $\ell_2$  instance-optimal bound might indicate. The reason is that, for any  $\Phi$ , such bounds consider the *worst case* signal over all possible  $x \in \mathbb{R}^N$ . Fortunately, this worst case is not typical. As a result, it is possible to derive much stronger results that consider any given signal  $x \in \mathbb{R}^N$  and establish that for most random  $\Phi$ , the recovery error of that signal  $x$  will be small.

## 5 Probabilistic instance-optimal bounds in $\ell_2$

For a given measurement operator  $\Phi$ , our bound in Theorem 4 applies uniformly to any signal in  $\mathbb{R}^N$ . However, a much sharper bound can be obtained by relaxing the deterministic requirement.

### 5.1 Signal recovery

Our first bound applies to the signal recovery problem, and we include the consideration of additive noise in the measurements.

**Theorem 5** *Suppose  $x \in \mathbb{R}^N$ . Let  $\mathcal{M}$  be a compact  $K$ -dimensional Riemannian submanifold of  $\mathbb{R}^N$  having condition number  $1/\tau$ , volume  $V$ , and geodesic covering regularity  $R$ . Fix  $0 < \epsilon < 1$  and  $0 < \rho < 1$ . Let  $\Phi$  be a random  $M \times N$  orthoprojector, chosen independently of  $x$ , with*

$$M = O\left(\frac{K \log(NV R \tau^{-1} \epsilon^{-1}) \log(1/\rho)}{\epsilon^2}\right). \quad (15)$$

*Let  $\eta \in \mathbb{R}^M$ , let  $y = \Phi x + \eta$ , and let the recovered estimate  $\hat{x}$  and the optimal estimate  $x^*$  be as defined in (10) and (11). If  $M \leq N$ , then with probability at least  $1 - \rho$  the following statement holds:*

$$\|x - \hat{x}\|_2 \leq (1 + 0.25\epsilon) \|x - x^*\|_2 + (2 + 0.32\epsilon) \|\eta\|_2 + \frac{\epsilon^2 \tau}{936N}. \quad (16)$$

**Proof:** See Appendix C.

The proof of this theorem, like that of Theorem 3, involves the Johnson-Lindenstrauss Lemma. Our proof of Theorem 5 extends the proof of Theorem 3 by adding the points  $x$  and  $x^*$  to the finite sampling of points drawn from  $\mathcal{M}$  that are used to establish (9).

Let us now compare and contrast our bound with the analogous results for sparsity-based CS. Like Theorem 2, we consider the problem of signal recovery in the presence of additive measurement noise. Both bounds relate the recovery error  $\|x - \hat{x}\|_2$  to the proximity of  $x$  to its nearest neighbor in the concise model class (either  $x_K$  or  $x^*$  depending on the model), and both bounds relate the recovery error  $\|x - \hat{x}\|_2$  to the amount  $\|\eta\|_2$  of additive measurement noise. However, Theorem 2 is a deterministic bound whereas Theorem 5 is probabilistic, and our bound (16) measures proximity to the concise model in the  $\ell_2$  norm, whereas (5) uses the  $\ell_1$  norm.

Our bound can also be compared with (6), as both are instance-optimal bounds in probability, and both use the  $\ell_2$  norm to measure proximity to the concise model. However, we note that unlike (6), our bound (16) allows the consideration of measurement noise.

Finally, we note that there is an additional term  $\frac{\epsilon^2 \tau}{936N}$  appearing on the right hand side of (16). This term becomes relevant only when both  $\|x - x^*\|_2$  and  $\|\eta\|_2$  are significantly smaller than the condition number  $\tau$ , since  $\epsilon^2 < 1$  and  $\frac{1}{936N} \ll 1$ . Indeed, in these regimes the signal recovery remains accurate (much smaller than  $\tau$ ), but the quantity  $\|x - \hat{x}\|_2$  may not remain strictly proportional to  $\|x - x^*\|_2$  and  $\|\eta\|_2$ . The bound may also be sharpened by artificially assuming a condition number  $1/\tau' > 1/\tau$  for the purpose of choosing a number of measurements  $M$  in (15). This will decrease the last term in (16) as  $\frac{\epsilon^2 \tau'}{936N}$ . In the case where  $\eta = 0$ , it is also possible to resort to the bound (14); this bound is inferior to (16) when  $\|x - x^*\|_2$  is large but ensures that  $\|x - \hat{x}\|_2 \rightarrow 0$  when  $\|x - x^*\|_2 \rightarrow 0$ .

## 5.2 Parameter estimation

Above we have derived a bound for the signal recovery problem, with an error metric that measures the discrepancy between the recovered signal  $\hat{x}$  and the original signal  $x$ .

However, in some applications it may be the case that the original signal  $x \approx x_{\theta^*}$ , where  $\theta^* \in \Theta$  is a parameter of interest. In this case we may be interested in using the compressive measurements  $y = \Phi x + \eta$  to solve the problem (12) and recover an estimate  $\hat{\theta}$  of the underlying parameter.

Of course, these two problems are closely related. However, we should emphasize that guaranteeing  $\|x - \hat{x}\|_2 \approx \|x - x^*\|_2$  does not automatically guarantee that  $d_{\mathcal{M}}(x_{\hat{\theta}}, x_{\theta^*})$  is small (and therefore does not ensure that  $d_{\Theta}(\hat{\theta}, \theta^*)$  is small). If the manifold is shaped like a horseshoe, for example, then it could be the case that  $x_{\theta^*}$  sits at the end of one arm but  $x_{\hat{\theta}}$  sits at the end of the opposing arm. These two points would be much closer in a Euclidean metric than in a geodesic one.

Consequently, in order to establish bounds relevant for parameter estimation, our concern focuses on guaranteeing that the geodesic distance  $d_{\mathcal{M}}(x_{\hat{\theta}}, x_{\theta^*})$  is itself small.

**Theorem 6** *Suppose  $x \in \mathbb{R}^N$ . Let  $\mathcal{M}$  be a compact  $K$ -dimensional Riemannian submanifold of  $\mathbb{R}^N$  having condition number  $1/\tau$ , volume  $V$ , and geodesic covering regularity  $R$ . Fix  $0 < \epsilon < 1$  and  $0 < \rho < 1$ . Let  $\Phi$  be a random  $M \times N$  orthoprojector, chosen independently of  $x$ , with*

$$M = O\left(\frac{K \log(NVR\tau^{-1}\epsilon^{-1}) \log(1/\rho)}{\epsilon^2}\right).$$

Let  $\eta \in \mathbb{R}^M$ , let  $y = \Phi x + \eta$ , and let the recovered estimate  $\hat{x}$  and the optimal estimate  $x^*$  be as defined in (10) and (11). If  $M \leq N$  and if  $1.16 \|\eta\|_2 + \|x - x^*\|_2 \leq \tau/5$ , then with probability at least  $1 - \rho$  the following statement holds:

$$d_{\mathcal{M}}(\hat{x}, x^*) \leq (4 + 0.5\epsilon) \|x - x^*\|_2 + (4 + 0.64\epsilon) \|\eta\|_2 + \frac{\epsilon^2 \tau}{468N}. \quad (17)$$

**Proof:** See Appendix D.

In several ways, this bound is similar to (16). Both bounds relate the recovery error to the proximity of  $x$  to its nearest neighbor  $x^*$  on the manifold and to the amount  $\|\eta\|_2$  of additive measurement noise. Both bounds also have an additive term on the right hand side that is small in relation to the condition number  $\tau$ .

In contrast, (17) guarantees that the recovered estimate  $\hat{x}$  is near to the optimal estimate  $x^*$  in terms of geodesic distance along the manifold. Establishing this condition required the additional assumption that  $1.16 \|\eta\|_2 + \|x - x^*\|_2 \leq \tau/5$ . Because  $\tau$  relates to the degree to which the manifold can curve back upon itself at long geodesic distances, this assumption prevents exactly the type of “horseshoe” problem that was mentioned above, where it may happen that  $d_{\mathcal{M}}(\hat{x}, x^*) \gg \|\hat{x} - x^*\|_2$ . Suppose, for example, it were to happen that  $\|x - x^*\|_2 \approx \tau$  and  $x$  was approximately equidistant from both ends of the horseshoe; a small distortion of distances under  $\Phi$  could then lead to an estimate  $\hat{x}$  for which  $\|x - \hat{x}\|_2 \approx \|x - x^*\|_2$  but  $d_{\mathcal{M}}(\hat{x}, x^*) \gg 0$ . Similarly, additive noise could cause a similar problem of “crossing over” in the measurement space. Although our bound provides no guarantee in these situations, we stress that under these circumstances, accurate parameter estimation would be difficult (or perhaps even unimportant) in the original signal space  $\mathbb{R}^N$ .

Finally, we revisit the situation where the original signal  $x \approx x_{\theta^*}$  for some  $\theta^* \in \Theta$  (with  $\theta^*$  satisfying (13)), where the measurements  $y = \Phi x + \eta$ , and where the recovered estimate  $\hat{\theta}$  satisfies (12). We consider the question of whether (17) can be translated into a bound on  $d_{\Theta}(\hat{\theta}, \theta^*)$ . As described in Section 3.1, in signal models where  $\mathcal{M}$  is isometric to  $\Theta$ , this is automatic: we have simply that

$$d_{\mathcal{M}}(x_{\hat{\theta}}, x_{\theta^*}) = d_{\Theta}(\hat{\theta}, \theta^*).$$

Such signal models are not nonexistent. Work by Donoho and Grimes [19], for example, has characterized a variety of articulated image classes for which (7) holds or for which  $d_{\mathcal{M}}(x_{\theta_1}, x_{\theta_2}) = C_8 d_{\Theta}(\theta_1, \theta_2)$  for some constant  $C_8 > 0$ . In other models it may hold that

$$C_9 d_{\mathcal{M}}(x_{\theta_1}, x_{\theta_2}) \leq d_{\Theta}(\theta_1, \theta_2) \leq C_{10} d_{\mathcal{M}}(x_{\theta_1}, x_{\theta_2})$$

for constants  $C_9, C_{10} > 0$ . Each of these relationships may be incorporated to the bound (17).

## 6 Conclusions and future work

In this paper, we have considered the tasks of signal recovery and parameter estimation using compressive measurements of a manifold-modeled signal. Although these problems differ substantially from the mechanics of sparsity-based signal recovery, we have seen a number of similarities that arise due to the low-dimensional geometry of the each of the concise models. First, we have seen that a sufficient number of compressive measurements can guarantee a stable embedding of either type of signal family, and the requisite number of measurements scales linearly with the information level of the signal. Second, we have seen that deterministic instance-optimal bounds in  $\ell_2$

are necessarily weak for both problems. Third, we have seen that probabilistic instance-optimal bounds in  $\ell_2$  can be derived that give the optimal scaling with respect to the signal proximity to the concise model and with respect to the amount of measurement noise. Thus, our work supports the growing empirical evidence that manifold-based models can be used with high accuracy in compressive signal processing.

As discussed in Section 3.3, there remain several active topics of research. One matter concerns the problem of non-differentiable manifolds that arise from certain classes of articulated image models. Based on preliminary and empirical work [20, 36], we believe that a combined multiscale regularization/measurement process is appropriate for such problems. However, a suitable theory should be developed to support this. A second topic of active research concerns fast algorithms for solving problems such as (10) and (12). Most successful approaches to date have combined initial coarse-scale discrete searches with iterative Newton-like refinements. Due to the problem-specific nuances that can arise in manifold models, it is unlikely that a single general-purpose algorithm analogous to  $\ell_1$ -minimization will emerge for solving these problems. Nonetheless, advances in these directions will likely be made by considering existing techniques for solving (11) and (13) in the native space, and perhaps by considering the multiscale measurement processes described above.

Finally, while we have not considered stochastic models for the parameter  $\theta$  or the noise  $\eta$ , it would be interesting to consider these situations as well. A starting point for such statistical analysis may be the constrained Cramér-Rao Bound formulations [30, 34] in which an unknown parameter is constrained to live along a low-dimensional manifold. However, the appropriate approach may once again be problem-dependent, as the nearest-neighbor estimators (12), (13) we describe can be biased for nonlinear or non-isometric manifolds.

## Acknowledgements

The author gratefully acknowledges Rice University, Caltech, and the University of Michigan, where he resided during portions of this research. An early version of Theorem 4 appeared in the author's Ph.D. thesis [36], under the supervision of Richard Baraniuk. Thanks to Rich and to the Rice CS research team for many stimulating discussions.

## A Geodesic covering regularity

We briefly review the definition of geodesic covering regularity and refer the reader to [3] for a deeper discussion.

**Definition 2** *Let  $\mathcal{M}$  be a compact Riemannian submanifold of  $\mathbb{R}^N$ . Given  $T > 0$ , the geodesic covering number  $G(T)$  of  $\mathcal{M}$  is defined as the smallest number such that there exists a set  $A$  of points on  $\mathcal{M}$ ,  $\#A = G(T)$ , so that for all  $x \in \mathcal{M}$ ,*

$$\min_{a \in A} d_{\mathcal{M}}(x, a) \leq T.$$

**Definition 3** *Let  $\mathcal{M}$  be a compact  $K$ -dimensional Riemannian submanifold of  $\mathbb{R}^N$  having volume  $V$ . We say that  $\mathcal{M}$  has geodesic covering regularity  $R$  for resolutions  $T \leq T_0$  if*

$$G(T) \leq \frac{R^K V K^{K/2}}{T^K} \tag{18}$$

for all  $0 < T \leq T_0$ .

## B Proof of Theorem 4

Fix  $\alpha \in [1 - \epsilon, 1 + \epsilon]$ . We consider any two points in  $w_a, w_b \in \mathcal{M}$  such that

$$\frac{\|\Phi w_a - \Phi w_b\|_2}{\|w_a - w_b\|_2} = \alpha,$$

and supposing that  $x$  is closer to  $w_a$ , i.e.,

$$\|x - w_a\|_2 \leq \|x - w_b\|_2,$$

but  $\Phi x$  is closer to  $\Phi w_b$ , i.e.,

$$\|\Phi x - \Phi w_b\|_2 \leq \|\Phi x - \Phi w_a\|_2,$$

we seek the maximum value that

$$\frac{\|x - w_b\|_2}{\|x - w_a\|_2}$$

may take. In other words, we wish to bound the worst possible “mistake” (according to our error criterion) between two candidate points on the manifold whose distance is scaled by the factor  $\alpha$ .

This can be posed in the form of an optimization problem

$$\begin{aligned} \max_{x \in \mathbb{R}^N, w_a, w_b \in \mathcal{M}} \frac{\|x - w_b\|_2}{\|x - w_a\|_2} \quad \text{s.t.} \quad & \|x - w_a\|_2 \leq \|x - w_b\|_2, \\ & \|\Phi x - \Phi w_b\|_2 \leq \|\Phi x - \Phi w_a\|_2, \\ & \frac{\|\Phi w_a - \Phi w_b\|_2}{\|w_a - w_b\|_2} = \alpha. \end{aligned}$$

For simplicity, we may expand the constraint set to include all  $w_a, w_b \in \mathbb{R}^N$ ; the solution to this larger problem is an upper bound for the solution to the case where  $w_a, w_b \in \mathcal{M}$ .

The constraints and objective function now are invariant to adding a constant to all three variables or to a constant rescaling of all three. Hence, without loss of generality, we set  $w_a = \mathbf{0}$  and  $\|x\|_2 = 1$ . This leaves

$$\begin{aligned} \max_{x, w_b \in \mathbb{R}^N} \|x - w_b\|_2 \quad \text{s.t.} \quad & \|x\|_2 = 1, \\ & \|x - w_b\|_2 \geq 1, \\ & \|\Phi x - \Phi w_b\|_2 \leq \|\Phi x\|_2, \\ & \frac{\|\Phi w_b\|_2}{\|w_b\|_2} = \alpha. \end{aligned}$$

We may safely ignore the second constraint (because of its relation to the objective function), and we may also square the objective function (to be later undone).

We recall that  $\Phi = \sqrt{N/M} \Xi$ , where  $\Xi$  is an  $M \times N$  matrix having orthonormal rows. We let  $\Xi'$  be an  $(N - M) \times N$  matrix having orthonormal rows that are orthogonal to the rows of  $\Xi$ , and we define  $\Phi' = \sqrt{N/M} \Xi'$ . It follows that for any  $x' \in \mathbb{R}^N$ ,

$$\|\Phi x'\|_2^2 + \|\Phi' x'\|_2^2 = (N/M) \|x'\|_2^2.$$

This leads to

$$\max_{x, w_b \in \mathbb{R}^N} (M/N) (\|\Phi x - \Phi w_b\|_2^2 + \|\Phi' x - \Phi' w_b\|_2^2)$$



subject to

$$\begin{aligned}\|\Phi x\|_2^2 + \|\Phi' x\|_2^2 &= N/M, \\ \|\Phi x - \Phi w_b\|_2^2 &\leq \|\Phi x\|_2^2, \\ \frac{\|\Phi w_b\|_2^2}{\|\Phi w_b\|_2^2 + \|\Phi' w_b\|_2^2} &= (M/N)\alpha^2.\end{aligned}$$

The last constraint may be rewritten as

$$\|\Phi' w_b\|_2^2 = \|\Phi w_b\|_2^2 \left( \frac{N}{M} \frac{1}{\alpha^2} - 1 \right).$$

We note that the  $\Phi$  and  $\Phi'$  components of each vector may be optimized separately (subject to the listed constraints) because they are orthogonal components of that vector. Define  $\beta$  to be the value of  $\|\Phi' w_b\|_2$  taken for the optimal solution  $w_b$ . We note that the constraints refer to the norm of the vector  $\Phi' w_b$  but not its direction. To maximize the objective function, then,  $\Phi' w_b$  must be parallel to  $\Phi' x$  but with the opposite sign. Equivalently, it must follow that

$$\Phi' w_b = -\beta \cdot \frac{\Phi' x}{\|\Phi' x\|_2}. \quad (19)$$

We now consider the second term in the objective function. From (19), it follows that

$$\begin{aligned}\|\Phi' x - \Phi' w_b\|_2^2 &= \left\| \Phi' x \left( 1 + \frac{\beta}{\|\Phi' x\|_2} \right) \right\|_2^2 \\ &= \|\Phi' x\|_2^2 \cdot \left( 1 + \frac{\beta}{\|\Phi' x\|_2} \right)^2.\end{aligned} \quad (20)$$

The third constraint also demands that

$$\beta^2 = \|\Phi w_b\|_2^2 \left( \frac{N}{M} \frac{1}{\alpha^2} - 1 \right).$$

Substituting into (20), we have

$$\begin{aligned}\|\Phi' x - \Phi' w_b\|_2^2 &= \|\Phi' x\|_2^2 \cdot \left( 1 + 2 \frac{\beta}{\|\Phi' x\|_2} + \frac{\beta^2}{\|\Phi' x\|_2^2} \right) \\ &= \|\Phi' x\|_2^2 + 2 \|\Phi' x\|_2 \|\Phi w_b\|_2 \sqrt{\frac{N}{M} \frac{1}{\alpha^2} - 1} \\ &\quad + \|\Phi w_b\|_2^2 \left( \frac{N}{M} \frac{1}{\alpha^2} - 1 \right).\end{aligned}$$

This is an increasing function of  $\|\Phi w_b\|_2$ , and so we seek the maximum value that  $\|\Phi w_b\|_2$  may take subject to the constraints. From the second constraint we see that  $\|\Phi x - \Phi w_b\|_2^2 \leq \|\Phi x\|_2^2$ ; thus,  $\|\Phi w_b\|_2$  is maximized by letting  $\Phi w_b = 2\Phi x$ . With such a choice of  $\Phi w_b$  we then have

$$\|\Phi x - \Phi w_b\|_2^2 = \|\Phi x\|_2^2.$$

We note that this choice of  $\Phi w_b$  *also* maximizes the first term of the objective function subject to the constraints.

We may now rewrite the optimization problem, in light of the above restrictions:

$$\begin{aligned} \max_{\Phi x, \Phi' x} (M/N) & \left( \|\Phi x\|_2^2 + \|\Phi' x\|_2^2 + 4 \|\Phi x\|_2 \|\Phi' x\|_2 \sqrt{\frac{N}{M} \frac{1}{\alpha^2} - 1} + 4 \|\Phi x\|_2^2 \left( \frac{N}{M} \frac{1}{\alpha^2} - 1 \right) \right) \\ \text{s.t. } & \|\Phi x\|_2^2 + \|\Phi' x\|_2^2 = \frac{N}{M}. \end{aligned}$$

We now seek to bound the maximum value that the objective function may take. We note that the single constraint implies that

$$\|\Phi x\|_2 \|\Phi' x\|_2 \leq \frac{1}{2} \left( \frac{N}{M} \right)$$

and that  $\|\Phi x\|_2 \leq \sqrt{N/M}$  (but because these cannot be simultaneously met with equality, our bound will not be tight). It follows that

$$\begin{aligned} & (M/N) \left( \|\Phi x\|_2^2 + \|\Phi' x\|_2^2 + 4 \|\Phi x\|_2 \|\Phi' x\|_2 \sqrt{\frac{N}{M} \frac{1}{\alpha^2} - 1} + 4 \|\Phi x\|_2^2 \left( \frac{N}{M} \frac{1}{\alpha^2} - 1 \right) \right) \\ & \leq (M/N) \left( \frac{N}{M} + 2 \frac{N}{M} \sqrt{\frac{N}{M} \frac{1}{\alpha^2} - 1} + 4 \frac{N}{M} \left( \frac{N}{M} \frac{1}{\alpha^2} - 1 \right) \right) \\ & = \frac{N}{M} \frac{4}{\alpha^2} - 3 + 2 \sqrt{\frac{N}{M} \frac{1}{\alpha^2} - 1}. \end{aligned}$$

Returning to the original optimization problem (for which we must now take a square root), this implies that

$$\frac{\|x - w_b\|_2}{\|x - w_a\|_2} \leq \sqrt{\frac{N}{M} \frac{4}{\alpha^2} - 3 + 2 \sqrt{\frac{N}{M} \frac{1}{\alpha^2} - 1}}$$

for any observation  $x$  that could be mistakenly paired with  $w_b$  instead of  $w_a$  (under a projection that scales the distance  $\|w_a - w_b\|_2$  by  $\alpha$ ). Considering the range of possible  $\alpha$ , the worst case may happen when  $\alpha = (1 - \epsilon)$ .  $\square$

## C Proof of Theorem 5

Following the proof of Theorem 3 (see [3]), we let  $\epsilon_1 = \frac{1}{13}\epsilon$  and  $T = \frac{\epsilon^2 \tau}{3100N}$ . We let  $A$  be a minimal set of points on the manifold  $\mathcal{M}$  such that, for every  $x' \in \mathcal{M}$ ,

$$\min_{a \in A} d_{\mathcal{M}}(x', a) \leq T. \quad (21)$$

We call  $A$  the set of *anchor points*. From (18) we have that  $\#A \leq \frac{R^K V K^{K/2}}{T^K}$ . The proof also describes a finite set of points  $B \supset A$  and applies the Johnson-Lindenstrauss Lemma to this set to conclude that

$$(1 - \epsilon_1) \|b_1 - b_2\|_2 \leq \|\Phi b_1 - \Phi b_2\|_2 \leq (1 + \epsilon_2) \|b_1 - b_2\|_2 \quad (22)$$

holds for all  $b_1, b_2 \in B$ . The cardinality of the set  $B$  dictates the requisite number of measurements  $M$  in (8).

For our purposes, we define a new set  $B' := B \cup \{x\} \cup \{x^*\}$ . Noting that  $\Phi$  is independent of both  $x$  and  $x^*$ , we may apply the Johnson-Lindenstrauss Lemma to  $B'$  instead and conclude that

(22) holds for all  $b_1, b_2 \in B'$ . This new set has cardinality  $|B'| \leq |B| + 2$ , and one may check that this does not change the order of the number of measurements required in (8).

Let  $\hat{a}$  denote the anchor point nearest to  $\hat{x}$  in terms of  $\ell_2$  distance in  $\mathbb{R}^N$ . It follows that  $\|\hat{x} - \hat{a}\|_2 \leq T$ . Since  $x, x^*, \hat{a} \in B'$ , we know that

$$(1 - \epsilon_1) \|x - \hat{a}\|_2 \leq \|\Phi x - \Phi \hat{a}\|_2 \leq (1 + \epsilon_1) \|x - \hat{a}\|_2$$

and

$$(1 - \epsilon_1) \|x - x^*\|_2 \leq \|\Phi x - \Phi x^*\|_2 \leq (1 + \epsilon_1) \|x - x^*\|_2.$$

Also, since  $\hat{x}, \hat{a} \in \mathcal{M}$ , we have from the conclusion of Theorem 3 that

$$(1 - \epsilon) \|\hat{x} - \hat{a}\|_2 \leq \|\Phi \hat{x} - \Phi \hat{a}\|_2 \leq (1 + \epsilon) \|\hat{x} - \hat{a}\|_2.$$

Finally, notice that by definition

$$\|x - x^*\|_2 \leq \|x - \hat{x}\|_2$$

and

$$\|(\Phi x + \eta) - \Phi \hat{x}\|_2 \leq \|(\Phi x + \eta) - \Phi x^*\|_2.$$

Now, combining all of these bounds and using several applications of the triangle inequality we have

$$\begin{aligned} \|x - \hat{x}\|_2 &\leq \|x - \hat{a}\|_2 + \|\hat{x} - \hat{a}\|_2 \\ &\leq \|x - \hat{a}\|_2 + T \\ &\leq \frac{1}{1 - \epsilon_1} \|\Phi x - \Phi \hat{a}\|_2 + T \\ &\leq \frac{1}{1 - \epsilon_1} (\|\Phi x - \Phi \hat{x}\|_2 + \|\Phi \hat{x} - \Phi \hat{a}\|_2) + T \\ &\leq \frac{1}{1 - \epsilon_1} (\|\Phi x - \Phi \hat{x} + \eta\|_2 + \|\eta\|_2 + \|\Phi \hat{x} - \Phi \hat{a}\|_2) + T \\ &\leq \frac{1}{1 - \epsilon_1} (\|\Phi x - \Phi x^* + \eta\|_2 + \|\eta\|_2 + \|\Phi \hat{x} - \Phi \hat{a}\|_2) + T \\ &\leq \frac{1}{1 - \epsilon_1} (\|\Phi x - \Phi x^*\|_2 + 2\|\eta\|_2 + \|\Phi \hat{x} - \Phi \hat{a}\|_2) + T \\ &\leq \frac{1}{1 - \epsilon_1} ((1 + \epsilon_1) \|x - x^*\|_2 + 2\|\eta\|_2 + \|\Phi \hat{x} - \Phi \hat{a}\|_2) + T \\ &\leq \frac{1}{1 - \epsilon_1} ((1 + \epsilon_1) \|x - x^*\|_2 + 2\|\eta\|_2 + T(1 + \epsilon)) + T \\ &= \frac{2\|\eta\|_2}{1 - \epsilon_1} + \frac{1 + \epsilon_1}{1 - \epsilon_1} \|x - x^*\|_2 + T \left( \frac{1 + \epsilon}{1 - \epsilon_1} + 1 \right). \end{aligned}$$

One can check that

$$\frac{1}{1 - \epsilon_1} \leq 1 + 0.16\epsilon,$$

$$\frac{1 + \epsilon_1}{1 - \epsilon_1} \leq 1 + 0.25\epsilon,$$

and

$$\frac{1 + \epsilon}{1 - \epsilon_1} \leq 1 + 1.31\epsilon.$$

Therefore,

$$\|x - \hat{x}\|_2 \leq (2 + 0.32\epsilon) \|\eta\|_2 + (1 + 0.25\epsilon) \|x - x^*\|_2 + \frac{\epsilon^2 \tau}{936N}.$$

□

## D Proof of Theorem 6

Using a simple triangle inequality and (16), we have

$$\|\hat{x} - x^*\|_2 \leq \|x - \hat{x}\|_2 + \|x - x^*\|_2 \leq (2 + 0.32\epsilon) \|\eta\|_2 + (2 + 0.25\epsilon) \|x - x^*\|_2 + \frac{\epsilon^2 \tau}{936N}. \quad (23)$$

Now, since both  $\hat{x}$  and  $x^*$  belong to  $\mathcal{M}$ , we can invoke Lemma 2.3 from [33], which states that if  $\|\hat{x} - x^*\|_2 \leq \tau/2$ , then

$$d_{\mathcal{M}}(\hat{x}, x^*) \leq \tau - \tau \sqrt{1 - 2\|\hat{x} - x^*\|_2 / \tau}. \quad (24)$$

(This lemma guarantees that two points separated by a small Euclidean distance are also separated by a small geodesic distance, and so the manifold does not “curve back” upon itself.) To apply this lemma, it is sufficient to know that

$$(2 + 0.32\epsilon) \|\eta\|_2 + (2 + 0.25\epsilon) \|x - x^*\|_2 + \frac{\epsilon^2 \tau}{936N} \leq \tau/2,$$

i.e., that

$$\frac{2 + 0.32\epsilon}{2 + 0.25\epsilon} \|\eta\|_2 + \|x - x^*\|_2 \leq \tau \left( \frac{\frac{1}{2} - \frac{\epsilon^2}{936N}}{2 + 0.25\epsilon} \right).$$

For the sake of neatness, we may tighten this condition to  $1.16\|\eta\|_2 + \|x - x^*\|_2 \leq \tau/5$ , which implies the sufficient condition above (since  $\epsilon < 1$ ). Thus, if  $\|x - x^*\|_2$  and  $\|\eta\|_2$  are sufficiently small (on the order of the condition number  $\tau$ ), then we may combine (23) and (24), giving

$$\begin{aligned} d_{\mathcal{M}}(\hat{x}, x^*) &\leq \tau - \tau \sqrt{1 - \frac{2}{\tau} \left( (2 + 0.32\epsilon) \|\eta\|_2 + (2 + 0.25\epsilon) \|x - x^*\|_2 + \frac{\epsilon^2 \tau}{936N} \right)} \\ &= \tau - \tau \sqrt{1 - \left( \frac{(4 + 0.64\epsilon)}{\tau} \|\eta\|_2 + \frac{(4 + 0.5\epsilon)}{\tau} \|x - x^*\|_2 + \frac{\epsilon^2}{468N} \right)}. \end{aligned} \quad (25)$$

Under the assumption that  $1.16\|\eta\|_2 + \|x - x^*\|_2 \leq \tau/5$ , it follows that

$$0 < \frac{(4 + 0.64\epsilon)}{\tau} \|\eta\|_2 + \frac{(4 + 0.5\epsilon)}{\tau} \|x - x^*\|_2 + \frac{\epsilon^2}{468N} < 1$$

and so

$$\begin{aligned} &\sqrt{1 - \left( \frac{(4 + 0.64\epsilon)}{\tau} \|\eta\|_2 + \frac{(4 + 0.5\epsilon)}{\tau} \|x - x^*\|_2 + \frac{\epsilon^2}{468N} \right)} \\ &> 1 - \left( \frac{(4 + 0.64\epsilon)}{\tau} \|\eta\|_2 + \frac{(4 + 0.5\epsilon)}{\tau} \|x - x^*\|_2 + \frac{\epsilon^2}{468N} \right). \end{aligned}$$

This allows us to simplify (25) and gives our final result: If  $1.16\|\eta\|_2 + \|x - x^*\|_2 \leq \tau/5$ , then

$$d_{\mathcal{M}}(\hat{x}, x^*) \leq (4 + 0.64\epsilon) \|\eta\|_2 + (4 + 0.5\epsilon) \|x - x^*\|_2 + \frac{\epsilon^2 \tau}{468N}.$$

□

## References

- [1] D. Achlioptas. Database-friendly random projections. In *Proc. Symp. on Principles of Database Systems (PODS '01)*, pages 274–281. ACM Press, 2001.
- [2] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constr. Approx.*, 2007. To appear.
- [3] R. G. Baraniuk and M. B. Wakin. Random projections of smooth manifolds. *Foundations of Computational Mathematics*, 2008. To appear.
- [4] D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk. Distributed compressed sensing. 2005. Preprint.
- [5] D. S. Broomhead and M. J. Kirby. The Whitney Reduction Network: A method for computing autoassociative graphs. *Neural Comput.*, 13(11):2595–2616, November 2001.
- [6] 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, February 2006.
- [7] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, August 2006.
- [8] E. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51(12), December 2005.
- [9] E. Candès and T. Tao. Decoding via linear programming. *IEEE Trans. Inform. Theory*, 51(12):4203–4215, December 2005.
- [10] E. Candès and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5425, December 2006.
- [11] E. J. Candès and M. B. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, 2008.
- [12] E.J. Candès. The restricted isometry property and its implications for compressed sensing. *Compte Rendus de l'Academie des Sciences, Paris*, 346:589–592, 2008.
- [13] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best  $k$ -term approximation. 2006. Preprint.
- [14] S. Dasgupta and A. Gupta. An elementary proof of the Johnson-Lindenstrauss lemma. Technical Report TR-99-006, Berkeley, CA, 1999.
- [15] M.A. Davenport, M.F. Duarte, M.B. Wakin, J.N. Laska, D. Takhar, K.F. Kelly, and R.G. Baraniuk. The smashed filter for compressive classification and target recognition. In *Proc. Computational Imaging V at SPIE Electronic Imaging*, January 2007.
- [16] R. DeVore, G. Petrova, and P. Wojtaszczyk. Instance-optimality in probability with an ell-1 decoder. 2008. Preprint.
- [17] D. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4), April 2006.
- [18] D. Donoho and Y. Tsaig. Extensions of compressed sensing. *Signal Processing*, 86(3):533–548, March 2006.
- [19] D. L. Donoho and C. Grimes. Image manifolds which are isometric to Euclidean space. *J. Math. Imaging Comp. Vision*, 23(1):5–24, July 2005.

- [20] M. Duarte, M. Davenport, M. Wakin, J. Laska, D. Takhar, K. Kelly, and R. Baraniuk. Multiscale random projections for compressive classification. In *Proc. IEEE Conf. on Image Processing (ICIP)*, September 2007.
- [21] M. F. Duarte, M. A. Davenport, D. Takbar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):83–91, 2008.
- [22] A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss. Improved Time Bounds for Near-Optimal Sparse Fourier Representations. In *Proc. SPIE Wavelets XI*, 2005.
- [23] A.C. Gilbert, M.J. Strauss, J.A. Tropp, and R. Vershynin. One sketch for all: fast algorithms for compressed sensing. In *Proc. ACM Symposium on Theory of Computing*, 2007.
- [24] D. Healy and D. J. Brady. Compression at the physical interface. *IEEE Signal Processing Magazine*, 25(2):67–71, 2008.
- [25] C. Hegde, M.B. Wakin, and R.G. Baraniuk. Random projections for manifold learning. In *Proc. Neural Information Processing Systems (NIPS)*, December 2007.
- [26] G. E. Hinton, P. Dayan, and M. Revow. Modeling the manifolds of images of handwritten digits. *IEEE Trans. Neural Networks*, 8(1):65–74, January 1997.
- [27] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. Symp. Theory of Computing*, pages 604–613, 1998.
- [28] 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 Workshop (DCAS)*, Dallas, Texas, October 2006.
- [29] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing MRI. *IEEE Signal Processing Magazine*, 25(2):72–82, 2008.
- [30] T. J. Moore, R. J. Kozick, and B. M. Sadler. The constrained Cramér–Rao bound from the perspective of fitting a model. *IEEE Signal Processing Letters*, 14(8):564–567, 2007.
- [31] D. Needell and J. A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. 2008. Preprint.
- [32] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. 2007. Preprint.
- [33] P. Niyogi, S. Smale, and S. Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete and Computational Geometry*, 39(1):419–441, 2008.
- [34] P. Stoica and B. C. Ng. On the Cramér–Rao bound under parametric constraints. *IEEE Signal Processing Letters*, 5(7):177–179, 1998.
- [35] M. Turk and A. Pentland. Eigenfaces for recognition. *J. Cogn. Neurosci.*, 3(1):71–83, 1991.
- [36] M. B. Wakin. *The Geometry of Low-Dimensional Signal Models*. PhD thesis, Department of Electrical and Computer Engineering, Rice University, Houston, TX, 2006.
- [37] M. B. Wakin, D. L. Donoho, H. Choi, and R. G. Baraniuk. The multiscale structure of non-differentiable image manifolds. In *Proc. Wavelets XI at SPIE Optics and Photonics*, San Diego, California, August 2005.