

RANDOM PROJECTIONS OF SIGNAL MANIFOLDS

Michael B. Wakin and Richard G. Baraniuk

Department of Electrical and Computer Engineering, Rice University

ABSTRACT

Random projections have recently found a surprising niche in signal processing. The key revelation is that the relevant *structure* in a signal can be preserved when that signal is projected onto a small number of random basis functions. Recent work has exploited this fact under the rubric of Compressed Sensing (CS): signals that are sparse in some basis can be recovered from small numbers of random linear projections. In many cases, however, we may have a more specific low-dimensional model for signals in which the signal class forms a nonlinear *manifold* in \mathbb{R}^N . This paper provides preliminary theoretical and experimental evidence that manifold-based signal structure can be preserved using small numbers of random projections. The key theoretical motivation comes from Whitney's Embedding Theorem, which states that a K -dimensional manifold can be embedded in \mathbb{R}^{2K+1} . We examine the potential applications of this fact. In particular, we consider the task of recovering a manifold-modeled signal from a small number of random projections. Thanks to our more specific model, we can recover certain signals using far fewer measurements than would be required using sparsity-driven CS techniques.

1. INTRODUCTION

Random projections have recently emerged as a surprisingly useful tool in signal processing. The key revelation is that the relevant *structure* in a signal can be preserved when that signal is projected onto a small number of random basis functions. Indeed, although some information is lost through such a projection, that information tends to be *incoherent* with the relevant structure in the signal. This fact (which can be formalized in various ways — see below) is useful for several reasons. For example, the process of acquiring and compressing a signal can be greatly simplified. In fact, this encoding process can proceed without knowledge of the structure that makes the signal compressible — in this sense random projections are a *universal* measurement tool. Another benefit is that random projections provide *dimensionality reduction*, which can significantly simplify certain computations.

One very general application of random projections deals with a cloud of P points in \mathbb{R}^N . The Johnson-Lindenstrauss Lemma [1] establishes that using $O(\log P)$ random projections, one can embed these points with minimal distortion of their pairwise distances (the “structure” that is preserved). This result is particularly useful for solving the *Nearest Neighbor* problem in computer science [2]. A primary benefit is that computation can be substantially reduced by operating in this lower-dimensional space.

Another, more recent field known as Compressed Sensing (CS) employs random projections in a more specific setting.




Fig. 1. Examples of parameterized signal models: (a) A wedgelet [13] is a parameterized edge on a square image block. (b) Arbitrary view of a simple object; as the viewpoint changes, the images trace out a nonlinear, non-differentiable manifold [14].

Rather than dealing with arbitrary clouds of points, CS examines classes of signals having sparse representations in some basis. For an N -sample signal that is K -sparse,¹ only $O(K \log N)$ random projections of the signal are required to reconstruct the signal with high probability. A variety of algorithms have been proposed for signal recovery [3–6], each requiring a slightly different number of projections. CS has many promising applications in signal acquisition, compression, medical imaging, and sensor networks [3, 4, 7–12].

CS employs a specific model (sparsity) in order to distinguish among an *infinite* number of points (which comprise unions of linear subspaces — see Section 3.1) based on their random projections. The principle of sparsity is applicable to a variety of signal types and is central to algorithms in denoising and compression. However it is still somewhat *generic*: in many cases we may have a more specific low-dimensional model for signals. Simple examples, illustrated in Figure 1, include straight edges in images (which can be parameterized by a slope and an offset) or multiple views of a fixed object (which can be parameterized by the camera position). In these cases, the signal class forms a nonlinear *manifold* in \mathbb{R}^N .

This paper provides preliminary theoretical and experimental evidence that manifold-based signal structure can be preserved using small numbers of random projections. The key theoretical motivation comes from Whitney's Embedding Theorem (see Section 2.2), which states that a K -dimensional manifold can be embedded in \mathbb{R}^{2K+1} . We examine the potential applications of this fact (extending the techniques introduced in [15, 16]). In particular, we consider the task of recovering a manifold-modeled signal from a small number of random projections. Thanks to our more specific model, the ability to recover the signal can be far superior to sparsity-driven CS techniques.

This paper is organized as follows. Section 2 introduces manifolds as a useful model for low-dimensional signal structure and discusses the implications of Whitney's Embedding Theorem. Section 3 considers the specific problem of recovering a signal

This work was supported by NSF, ONR, AFOSR, and the Texas Instruments Leadership University Program.

Email: {wakin, richb}@rice.edu. Web: dsp.rice.edu/cs

¹By K -sparse, we mean that the signal can be written as a sum of K elements from some basis or dictionary in \mathbb{R}^N .

from its random projections. Section 4 presents an experimental application in image edge detection, and Section 5 concludes with a summary of ongoing research directions.

2. SIGNAL MANIFOLDS

2.1. Articulated signal models

In many cases where one has a low-dimensional notion of signal structure, the resulting signal class manifests itself as a nonlinear *submanifold* embedded in the signal space. Suppose, for example, that one knew (or could conceive of) a low-dimensional parameter $\theta \in \Theta$ that somehow controlled the generation of the signal. Examples of this very generic notion include: images of a straight edge (parameterized by 2 variables: a slope and an offset — see Figure 1(a)), multiple views of a fixed object (~ 5 parameters for the camera position — see Figure 1(b)), signals of unknown translation (1 parameter for shift — see Figure 2), or the output of some explicitly parameterized or articulated physical system [14, 17]. In each of these cases, if we denote by $f_\theta \in \mathbb{R}^N$ the signal formed with a particular parameter θ , then the corresponding family $\mathcal{M} = \{f_\theta : \theta \in \Theta\}$ forms a submanifold of \mathbb{R}^N . We let K denote the dimension of θ , which under mild assumptions matches the dimension of \mathcal{M} .

The above are just a few example scenarios in which manifolds may arise in signal processing. Recent investigations have examined the structural properties of particular signal manifolds, including the image manifolds in Figure 1 [14, 17]. One surprising finding is that manifolds generated by articulated images having sharp edges are *nowhere differentiable*; instead they have an inherent *multiscale structure* that can be characterized and exploited in image processing.

2.2. Random projections of manifolds

This paper exploits another possibly surprising fact: much of a signal manifold’s structure is actually preserved when it is projected from \mathbb{R}^N onto a random lower-dimensional subspace. The result follows from the proof of Whitney’s (Easy) Embedding Theorem.

Theorem 2.1 [18] *Let \mathcal{M} be a compact Hausdorff C^r K -dimensional manifold, with $2 \leq r \leq \infty$. Then there is a C^r embedding of \mathcal{M} in \mathbb{R}^{2K+1} .*

The proof of this theorem is highly insightful and considers the normalized secant set of the manifold

$$\Sigma = \left\{ \frac{x - x'}{\|x - x'\|_2} : x, x' \in \mathcal{M} \right\}.$$

Roughly speaking, the secant set forms a $2K$ -dimensional subset of the $(N - 1)$ -dimensional unit sphere (which equates with the space of projections from \mathbb{R}^N to \mathbb{R}^{N-1}), and so there exists a projection from \mathbb{R}^N to \mathbb{R}^{N-1} that projects \mathcal{M} injectively (without overlap). This can be repeated until reaching \mathbb{R}^{2K+1} . In signal processing, this secant set has been explicitly employed in order to find the optimal projection vectors for a given manifold (see [15, 16], which also provide interesting and insightful discussions).

This paper builds upon the following useful observation: Using identical arguments and assuming mild conditions on the signal manifold \mathcal{M} , it also follows that a *randomly* chosen projection

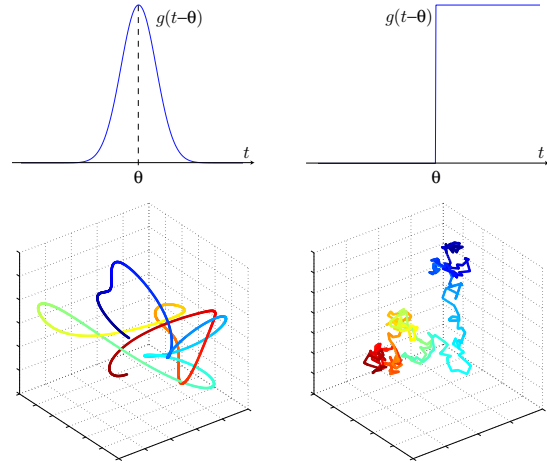


Fig. 2. Top row: The articulated signals $f_\theta(t) = g(t - \theta)$ are defined via shifts of a primitive function g , where g is (left) a Gaussian pulse or (right) a step function. The resulting signals trace out 1-D manifolds in \mathbb{R}^N . Bottom row: Projection of manifolds from \mathbb{R}^N onto 3 random functions; the color/shading corresponds to different values of $\theta \in \mathbb{R}$.

of the manifold from \mathbb{R}^N to \mathbb{R}^{2K+1} will be invertible with high probability.

To fix notation, we will consider signals $x \in \mathbb{R}^N$ and denote the M -dimensional random projection operator by \mathcal{P} . This operator can be implemented by constructing M random vectors in \mathbb{R}^N , which will span an M -dimensional linear subspace of \mathbb{R}^N , onto which x is orthogonally projected. (The computations simply involve inner products, and “signal processing” in this projected space involves only these M coefficients.) As an example, Figure 2 shows the random projection of two 1-dimensional (1-D) manifolds from \mathbb{R}^N to \mathbb{R}^3 and reveals that the particular characteristics of the manifold (such as differentiability) play a critical role. Although many interesting signal manifolds do not satisfy the criteria of Theorem 2.1, the theorem still provides a useful motivation; this paper provides empirical justification for scenarios outside of the Whitney criteria.

3. APPLICATION: SIGNAL RECOVERY

Random projections have many potential applications in manifold-based signal processing, including signal acquisition, sensor networks, and compression. In this paper, we focus on one fundamental task: reconstructing a signal x from its projection $\mathcal{P}x$. This is facilitated by assuming that x lives near the manifold \mathcal{M} and by exploiting the injectivity of the projection $\mathcal{P}\mathcal{M}$. The most relevant questions include: How can x be recovered from $\mathcal{P}x$? How many dimensions M are required for stability? How accurate will the reconstruction be? This paper provides preliminary theoretical insights into these topics and provides promising numerical experiments.

3.1. Recovering manifold-based signals

When the signal is “noiseless” — that is $x \in \mathcal{M}$ — then Whitney’s Embedding Theorem implies that x should be recoverable from $\mathcal{P}x$ (with high probability), assuming the dimension $M \geq 2K + 1$.

In some situations it may indeed be to further reduce M , though in general we expect $M \geq K$ to be a strict requirement.

Concrete algorithms for recovery are beyond the scope of this paper but essentially involve searching the projected manifold

$$\hat{x} = \arg \min_{x' \in \mathcal{M}} \|\mathcal{P}x - \mathcal{P}x'\|_2. \quad (1)$$

Figure 3 (discussed further in Section 4) shows a simple demonstration of recovery from $2K + 1$ projections.

Relation to Compressed Sensing: The $2K$ factor plays a very important role in CS, which can be reinterpreted in light of the manifold viewpoint. The signal “manifold” in the CS setting consists of a union of K -dimensional hyperplanes. The secant set for this turns out to be a union of $2K$ -dimensional hyperplanes (which loses 1 dimension after normalization). It follows that a random projection of a length- N , K -sparse signal onto $M = 2K$ dimensions is invertible with probability one [12]. (However, tractable recovery from $\mathcal{P}x$ requires slightly more measurements [3–6].)

3.2. Recovering near-manifold signals

A potentially more interesting scenario arises when the manifold is only an *approximation* for the signal class. Examples include edges that are not entirely straight or manifold-based signals corrupted by noise. We would like our recovery algorithm (1) to provide robust recovery of such signals. Ensuring such robustness now requires some notion of the *quality* of the manifold’s embedding into \mathbb{R}^M . Intuitively, if two far-away points $x, x' \in \mathcal{M}$ were to be mapped onto nearby points, then accurate recovery of any signals falling between x and x' would be difficult.

This notion can be made precise by defining [15, 16]

$$\kappa := \inf_{x, x' \in \mathcal{M}; x \neq x'} \frac{\|\mathcal{P}x - \mathcal{P}x'\|_2}{\|x - x'\|_2}. \quad (2)$$

Using this quantity, we can bound the error in a recovered signal relative to the original signal’s distance from the manifold. We have the following theorem, with proof omitted for brevity.

Theorem 3.1 *Let $x \in \mathbb{R}^N$ be an observation, and let \hat{x} be the estimation recovered from the projection $\mathcal{P}x$ via (1). Define $\delta = \inf_{x' \in \mathcal{M}} \|x - x'\|_2$, and let $\gamma = \|x - \hat{x}\|_2$. Then*

$$\frac{\gamma}{\delta} \leq \sqrt{\frac{4}{\kappa^2} - 3} + 2\sqrt{\frac{1}{\kappa^2} - 1}.$$

As $\kappa \rightarrow 1$, the bound on the right reduces simply to 1, and as $\kappa \rightarrow 0$, the bound grows as $2/\kappa$. We stress that this is a *worst case* bound and that the accuracy is often significantly better in practice (see Section 4). Moreover, we note that κ itself is a worst case bound relating to two points on the manifold. Many other pairwise distances may in fact be much more well preserved. (This suggests that there may exist more appropriate criteria for analysis — a topic of ongoing work.)

Another relevant topic of investigation is relating κ to M in order to discern the number of measurements required for robust signal recovery. Clearly, the anticipated κ increases to 1 as $M \rightarrow N$. However, arguments following Whitney’s Embedding Theorem imply only that $\kappa > 0$ with high probability when $M \geq 2K + 1$. A more rigorous theoretical investigation is currently underway




Fig. 3. Estimating image edge structure from a 256-pixel block. (a) Original 16×16 block. (b) Manifold-based recovery from 5 random projections. (c) Traditional CS recovery from 5 random projections using OMP algorithm [6] with Haar wavelets. (d) OMP recovery from 50 random projections. Perfect OMP recovery requires 70 or more random projections.

(which will relate the expected κ to the properties of the manifold such as curvature, volume, etc.); however, the next section provides numerical experiments to characterize the performance of our recovery algorithm. Finally, we mention that the algorithms introduced in [15, 16] aim specifically to find projection directions that maximize the quantity κ . However these lack the *universal* applicability of random projections.

4. EXAMPLE: EDGE DETECTION

In order to illustrate the basic principles in action, we now consider a simple image processing application. It is well known that much of the critical information in images is carried by the edges, which themselves have simple low-dimensional descriptions. One localized model for edge structure is provided by a *wedgelet* [13], a parameterized edge on a square image block (see Figure 1(a)).

We can consider the following task: given random projections of a local image segment, recover an approximation to the local edge structure. This can be formulated in our setting if we consider $x \in \mathbb{R}^N$ as the original image and $\mathcal{P}x$ as the random observation. We can then search the projected wedgelet manifold \mathcal{PM} for the closest match to the observation. While the wedgelet manifold is known not to be differentiable [14, 17], we consider it to be an interesting case and use this experiment to demonstrate the potential for applications beyond the specific assumptions of Theorem 2.1.

For the first experiment (Figure 3), we examine a perfect edge originating from a clean image $x \in \mathcal{M}$. We measure a 16×16 image block (256 pixels) using $2 \cdot 2 + 1 = 5$ random projections. To recover \hat{x} we use a sampling of the projected manifold and simply find the nearest neighbor to $\mathcal{P}x$. (Note that this experiment assumes the two grayscale values are known for the edge, or equivalently, that the mean and energy of x are provided along with the measurements.) While the sampling grid for the manifold search did not contain $\mathcal{P}x$ precisely, we see in Figure 3(b) that a very close approximation is recovered. In contrast, using traditional CS techniques to recover x from its random projections (by assuming that the image is sparse in the 2-D Haar wavelet domain) requires an order of magnitude more measurements. This comparison is a bit artificial but emphasizes how manifold models can be useful for recovering structured information from few random measurements.

For the second experiment (Figure 4) we analyze the robustness of the recovery process. For this we examine a 256×256 portion of the *Peppers* test image. We break the image into square blocks of size 16×16 , measure each one using 10 random projections (plus we include the mean and energy of each block), and then search the projected manifold to estimate a wedgelet on each block. We see from the figure that the recovery is fairly robust and accurately recovers most of the prominent edge structure, even though none of the original image blocks perfectly fits the




Fig. 4. (a) Original 256×256 Peppers image. (b) Wedgelet estimation on 16×16 pixel tiles, using 10 random projections (plus the mean and energy) on each tile, for a total of $(10 + 2) \cdot 256 = 3072$ measurements. (c) Best-possible wedgelet estimation, which would require all $256^2 = 65536$ pixel values. (d) Traditional CS-based recovery (from 3072 global random projections) using greedy pursuit to find a sparse approximation in the projected wavelet (D8) basis.

wedgelet model. The recovery is also fast, taking less than one second for the entire image. For point of comparison we include the best-possible wedgelet approximation to the image, which would require all 256 numbers per block to determine. In spite of the relatively small κ generated by the random projections (approximately 0.05 when computed using (2) over all pairs of wedgelets in our sampled grid), each wedgelet estimate is no more than 3 times worse than the best-possible wedgelet estimate (as measured by γ/δ in Theorem 3.1). For a second point of comparison with the wedgelet estimates, we also include the CS-based recovery of the whole image from an equivalent number of total measurements, using $(10 + 2) \cdot 256 = 3072$ global random projections. Though slightly better in terms of mean-square error, this approximation fails to prominently represent the edge structure (it also takes several minutes to compute using our software). We stress again, though, that the main purpose of the wedgelet experiment is to illustrate the robustness of recovery on natural image segments, some of which are not well-modeled using wedgelets.

5. CONCLUSIONS

This paper has provided theoretical and experimental evidence that manifold-based signal structure can be preserved using a small number of random projections. Thanks to this more specific model, we can recover certain signals using far fewer measurements than would be required using sparsity-driven CS techniques. Much work remains and is ongoing. For example, we wish to develop fast algorithms for signal recovery in the projected space (to deal with manifolds of dimension higher than two). We also aim to supplement existing techniques in CS, which may require methods for recovering local signal structure (such as wedgelets) from global image measurements. Conversely, this approach could motivate a somewhat localized CS measurement scheme. Finally, a more rigorous theoretical investigation is underway, particularly to quantify the necessary number of random measurements for robust signal recovery.

Acknowledgements: Thanks to Michael Kirby for helping to inspire this work and for enlightening discussions, and thanks to David Donoho, Hyeokho Choi, Venkat Chandrasekaran, Mark Davenport, Marco Duarte, Dror Baron, Shriram Sarvotham, and Robert Nowak for many helpful conversations.

6. REFERENCES

- [1] S. Dasgupta and A. Gupta, "An elementary proof of the Johnson-Lindenstrauss lemma," Tech. Rep. TR-99-006, Berkeley, CA, 1999.
- [2] P. Indyk and R. Motwani, "Approximate nearest neighbors: Towards removing the curse of dimensionality," in *Proc. 30th Ann. ACM Sympos. Theory Comput.*, 1998, pp. 604–613.
- [3] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," 2004, Preprint.
- [4] E. Candès and T. Tao, "Near optimal signal recovery from random projections and universal encoding strategies," 2004, Preprint.
- [5] D. Donoho, "Compressed sensing," 2004, Preprint.
- [6] J. Tropp and A. C. Gilbert, "Signal recovery from partial information via orthogonal matching pursuit," 2005, Preprint.
- [7] E. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n ," 2005, Preprint.
- [8] E. Candès and T. Tao, "Decoding by linear programming," 2004, Preprint.
- [9] D. Donoho and Y. Tsaig, "Extensions of compressed sensing," 2004, Preprint.
- [10] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," 2005, Preprint.
- [11] M. B. Wakin, M. F. Duarte, S. Sarvotham, D. Baron, and R. G. Baraniuk, "Recovery of jointly sparse signals from few random projections," in *Proc. Neural Inform. Processing Systems – NIPS*, 2005.
- [12] D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed compressed sensing," 2005, Preprint. Available at www.dsp.rice.edu/cs.
- [13] D. L. Donoho, "Wedgelets: Nearly-minimax estimation of edges," *Annals of Stat.*, vol. 27, pp. 859–897, 1999.
- [14] 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, August 2005.
- [15] D. S. Broomhead and M. Kirby, "A new approach for dimensionality reduction: Theory and algorithms," *SIAM J. of Applied Mathematics*, vol. 60, no. 6, 2000.
- [16] D. S. Broomhead and M. J. Kirby, "The Whitney Reduction Network: a method for computing autoassociative graphs," *Neural Computation*, vol. 13, pp. 2595–2616, 2001.
- [17] D. L. Donoho and C. Grimes, "Image manifolds isometric to Euclidean space," *J. Math. Imaging and Computer Vision*, 2003, To appear.
- [18] M. W. Hirsch, *Differential Topology*, vol. 33 of *Graduate Texts in Mathematics*, Springer, 1976.