

Compressive Sensing

Richard Baraniuk

Rice University

[Lecture Notes in *IEEE Signal Processing Magazine*]

Volume 24, July 2007

1 Scope

The Shannon/Nyquist sampling theorem tells us that in order to not lose information when uniformly sampling a signal we must sample at least two times faster than its bandwidth. In many applications, including digital image and video cameras, the Nyquist rate can be so high that we end up with too many samples and must compress in order to store or transmit them. In other applications, including imaging systems (medical scanners, radars) and high-speed analog-to-digital converters, increasing the sampling rate or density beyond the current state-of-the-art is very expensive.

In this lecture, we will learn about a new technique that tackles these issues using *compressive sensing* [1, 2]. We will replace the conventional sampling and reconstruction operations with a more general linear measurement scheme coupled with an optimization in order to acquire certain kinds of signals at a rate significantly below Nyquist.

2 Relevance

The ideas presented here can be used to illustrate the links between data acquisition, linear algebra, basis expansions, inverse problems, compression, dimensionality reduction, and optimization in a variety of courses, from undergraduate or graduate digital signal processing to statistics and applied mathematics.

3 Prerequisites

The prerequisites required for understanding and teaching this material are linear algebra, basic optimization, and basic probability.

4 Problem Statement

Nyquist-rate sampling completely describes a signal by exploiting its bandlimitedness. Our goal is to reduce the number of measurements required to completely describe a signal by exploiting its *compressibility*. The twist is that our measurements are not point samples but more general linear functionals of the signal.

Consider a real-valued, finite-length, one-dimensional, discrete-time signal \mathbf{x} , which we view as an $N \times 1$ column vector in \mathbb{R}^N with elements $x[n]$, $n = 1, 2, \dots, N$. We treat an image or higher-dimensional data by vectorizing it into a long one-dimensional vector.

Any signal in \mathbb{R}^N can be represented in terms of a basis of $N \times 1$ vectors $\{\boldsymbol{\psi}_i\}_{i=1}^N$. For simplicity, assume that the basis is orthonormal. Forming the $N \times N$ basis matrix $\boldsymbol{\Psi} := [\boldsymbol{\psi}_1 | \boldsymbol{\psi}_2 | \dots | \boldsymbol{\psi}_N]$ by stacking the vectors $\{\boldsymbol{\psi}_i\}$ as columns, we can express any signal \mathbf{x} as

$$\mathbf{x} = \sum_{i=1}^N s_i \boldsymbol{\psi}_i \quad \text{or} \quad \mathbf{x} = \boldsymbol{\Psi} \mathbf{s} \quad (1)$$

where \mathbf{s} is the $N \times 1$ column vector of weighting coefficients $s_i = \langle \mathbf{x}, \boldsymbol{\psi}_i \rangle = \boldsymbol{\psi}_i^T \mathbf{x}$ and where \cdot^T denotes the (Hermitian) transpose operation. Clearly, \mathbf{x} and \mathbf{s} are equivalent representations of the same signal, with \mathbf{x} in the time domain and \mathbf{s} in the $\boldsymbol{\Psi}$ domain.

We will focus on signals that have a *sparse* representation, where \mathbf{x} is a linear combination of just K basis vectors, with $K \ll N$. That is, only K of the s_i in (1) are nonzero and $(N - K)$ are zero. Sparsity is motivated by the fact that many natural and manmade signals are *compressible* in the sense that there exists a basis $\boldsymbol{\Psi}$ where the representation (1) has just a few large coefficients and many small coefficients. Compressible signals are well approximated by K -sparse representations; this is the basis of *transform coding* [3]. For example, natural images tend to be compressible in the discrete cosine transform (DCT) and wavelet bases [3] on which the JPEG and JPEG-2000 compression standards are based. Audio signals and many communication signals are compressible in a localized Fourier basis.

Transform coding plays a central role in data acquisition systems like digital cameras where the number of samples is high but the signals are compressible. In this framework, we acquire the full N -sample signal \mathbf{x} ; compute the complete set of transform coefficients $\{s_i\}$ via $\mathbf{s} = \boldsymbol{\Psi}^T \mathbf{x}$; locate the K largest coefficients and discard the $(N - K)$ smallest coefficients; and encode the K values and locations of the largest coefficients. (In practice, we also convert the values and locations to digital bits.)

Unfortunately, the sample-then-compress framework suffers from three inherent inefficiencies: First, we must start with a potentially large number of samples N even if the ultimate desired K is small. Second, the encoder must compute all of the N transform coefficients $\{s_i\}$, even though it will discard all but K of them. Third, the encoder faces the overhead of encoding the locations of the large coefficients.

As an alternative, we will study a more general data acquisition approach that condenses the signal directly into a compressed representation without going through the intermediate stage of taking N samples. Consider the more general *linear measurement process* that computes $M < N$




Figure 1: (a) Compressive sensing measurement process with (random Gaussian) measurement matrix Φ and discrete cosine transform (DCT) matrix Ψ . The coefficient vector \mathbf{s} is sparse with $K = 4$. (b) Measurement process in terms of the matrix product $\Theta = \Phi\Psi$ with the four columns corresponding to nonzero s_i highlighted. The measurement vector \mathbf{y} is a linear combination of these four columns.

inner products between \mathbf{x} and a collection of vectors $\{\phi_j\}_{j=1}^M$ as in $y_j = \langle \mathbf{x}, \phi_j \rangle$. Stacking the measurements y_j into the $M \times 1$ vector \mathbf{y} and the measurement vectors ϕ_j^T as rows into an $M \times N$ matrix Φ and substituting in (1), we can write

$$\mathbf{y} = \Phi\mathbf{x} = \Phi\Psi\mathbf{s} = \Theta\mathbf{s} \quad (2)$$

where $\Theta := \Phi\Psi$ is an $M \times N$ matrix. See Figure 1(a) for a pictorial depiction of (2). Note that the measurement process is non-adaptive; that is, Φ does not depend in any way on the signal \mathbf{x} .

Our goal in the following is to design a measurement matrix Φ and a reconstruction algorithm for K -sparse and compressible signals that require only $M \approx K$ or slightly more measurements, or about as many measurements as the number of coefficients encoded in a traditional transform coder. Our approach is based on the theory of *compressive sensing* introduced recently in [1, 2].

5 Solution

The solution consists of two steps. In the first step, we design a *stable measurement matrix* Φ that ensures that the salient information in any K -sparse or compressible signal is not damaged by the dimensionality reduction from $\mathbf{x} \in \mathbb{R}^N$ down to $\mathbf{y} \in \mathbb{R}^M$. In the second step, we develop a *reconstruction algorithm* to recover \mathbf{x} from the measurements \mathbf{y} . Initially, we focus on exactly K -sparse signals.

5.1 Stable measurement matrix

First, we design the measurement side of the data acquisition system, which is based around the matrix Φ . We aim to make M measurements (the vector \mathbf{y}) from which we can stably reconstruct the length- N signal \mathbf{x} , or equivalently its sparse coefficient vector \mathbf{s} in the basis Ψ . Clearly reconstruction will not be possible if the measurement process damages the information in \mathbf{x} . Unfortunately, this is the case in general: Since the measurement process is linear and defined in terms of the matrices Φ and Ψ , solving for \mathbf{s} given \mathbf{y} in (2) is just a linear algebra problem, and with

$M < N$, there are fewer equations than unknowns, making the solution ill-posed in general.

However, the K -sparsity of \mathbf{s} comes to our rescue. In this case the measurement vector \mathbf{y} is just a linear combination of the K columns of Θ whose corresponding $s_i \neq 0$ (see Figure 1(b)). Hence, if we knew a priori which K entries of \mathbf{s} were nonzero, then we could form an $M \times K$ system of linear equations to solve for these nonzero entries, where now the number of equations M equals or exceeds the number of unknowns K . A necessary and sufficient condition to ensure that this $M \times K$ system is well-conditioned — and hence sports a stable inverse — is that for any vector \mathbf{v} sharing the same K nonzero entries as \mathbf{s} we have

$$1 - \epsilon \leq \frac{\|\Theta \mathbf{v}\|_2}{\|\mathbf{v}\|_2} \leq 1 + \epsilon \quad (3)$$

for some $\epsilon > 0$. In words, the matrix Θ must preserve the lengths of these particular K -sparse vectors.

Of course, in practice we are not going to know the locations of the K nonzero entries in \mathbf{s} . Interestingly, one can show that a sufficient condition for a stable inverse for both K -sparse and compressible signals is for Θ to satisfy (3) for an arbitrary $3K$ -sparse vector \mathbf{v} . This is the so-called *restricted isometry property* (RIP) [1].

An alternative approach to stability is to ensure that the measurement matrix Φ is *incoherent* with the sparsifying basis Ψ in the sense that the vectors $\{\phi_j\}$ cannot sparsely represent the vectors $\{\psi_i\}$ and vice versa [1, 2, 4]. The classical example features delta spikes and Fourier sinusoids playing the roles of $\{\phi_j\}$ and $\{\psi_i\}$; the Fourier uncertainty principle immediately yields the incoherence.

So, given a sparsifying basis Ψ , how do we construct a measurement matrix Φ such that $\Theta = \Phi\Psi$ has the RIP? Unfortunately, merely verifying that a given Θ has the RIP is combinatorially complex; we must verify (3) for each of the $\binom{N}{K}$ possible combinations of K nonzero entries in the length- N vector \mathbf{v} .

In compressive sensing, we sidestep this issue by selecting Φ as a random matrix. For example, we draw the matrix elements $\phi_{j,i}$ as independent and identically distributed (iid) random variables from a zero-mean, $1/N$ -variance Gaussian density (white noise) [1, 2, 5]. Then, the measurements \mathbf{y} are merely M different randomly weighted linear combinations of the elements of \mathbf{x} (recall Figure 1(a) and note the random structure of Φ).

A Gaussian Φ has two interesting and useful properties. First, Φ is incoherent with the basis $\Psi = \mathbf{I}$ of delta spikes with high probability, since it takes fully N spikes to represent each row of Φ . More rigorously, using concentration of measure arguments, an $M \times N$ iid Gaussian matrix $\Theta = \Phi\mathbf{I} = \Phi$ can be shown to have the RIP with high probability if $M \geq cK \log(N/K)$, with c a small constant [1, 2, 5]. Therefore, we can expect to recover length- N , K -sparse and compressible

signals with high probability from just $M \geq cK \log(N/K) \ll N$ random Gaussian measurements. Second, thanks to the properties of the iid Gaussian distribution generating Φ , the matrix $\Theta = \Phi\Psi$ is also iid Gaussian regardless of the choice of (orthonormal) sparsifying basis matrix Ψ . Thus, random Gaussian measurements Φ are *universal* in the sense that $\Theta = \Phi\Psi$ has the RIP with high probability for every possible Ψ .

Among many others, Rademacher matrices with random ± 1 entries can also be shown to have the RIP and universality properties [5].

5.2 Signal reconstruction algorithms

The RIP provides the theoretical guarantee that a K -sparse or compressible signal can be fully described by the M measurements in \mathbf{y} , but it does not tell us how to recover it. The signal reconstruction step must take the measurements \mathbf{y} , the random measurement matrix Φ (or the random seed that generated it), and the sparsifying basis Ψ and regenerate the length- N signal \mathbf{x} , or equivalently its sparse coefficient vector \mathbf{s} . We again focus initially on K -sparse signals.

Since $M < N$ in (2), there are infinitely many \mathbf{s}' that satisfy $\Theta\mathbf{s}' = \mathbf{y}$; they all lie on the $(N - M)$ -dimensional hyperplane $\mathcal{H} := \mathcal{N}(\Theta) + \mathbf{s}$ in \mathbb{R}^N corresponding to the null space $\mathcal{N}(\Theta)$ of Θ translated to the true sparse solution \mathbf{s} . This is because if $\Theta\mathbf{s} = \mathbf{y}$ then $\Theta(\mathbf{s} + \mathbf{r}) = \mathbf{y}$ for any vector \mathbf{r} in the null space. Therefore, our goal is to find the signal's sparse coefficient vector \mathbf{s} in the translated null space.

Define the ℓ_p norm of the vector \mathbf{s} as $(\|\mathbf{s}\|_p)^p = \sum_{i=1}^N |s_i|^p$. When $p = 0$ we obtain the ℓ_0 “norm” that counts the number of non-zero entries in \mathbf{s} ; hence a K -sparse vector has ℓ_0 norm K .

Minimum ℓ_2 norm reconstruction: The classical approach to solving inverse problems of this type is by least squares; that is, we select the vector in the translated nullspace \mathcal{H} with smallest ℓ_2 norm (energy):

$$\hat{\mathbf{s}} = \operatorname{argmin} \|\mathbf{s}'\|_2 \quad \text{such that } \Theta\mathbf{s}' = \mathbf{y}. \quad (4)$$

There is even a convenient closed-form solution $\hat{\mathbf{s}} = \Theta^T(\Theta\Theta^T)^{-1}\mathbf{y}$. But unfortunately when the vector \mathbf{s} we seek is K -sparse, ℓ_2 minimization will almost never find it. What we obtain instead is a nonsparse $\hat{\mathbf{s}}$ with plenty of ringing (more on this below in Section 6.1).

Minimum ℓ_0 norm reconstruction: Since the ℓ_2 norm in (4) does not reflect signal sparsity, a logical alternative is to search for the sparsest vector in the translated null space \mathcal{H} :

$$\hat{\mathbf{s}} = \operatorname{argmin} \|\mathbf{s}'\|_0 \quad \text{such that } \Theta\mathbf{s}' = \mathbf{y}. \quad (5)$$

It can be shown that with just $M = K + 1$ iid Gaussian measurements, this optimization will recover a K -sparse signal exactly with high probability [6]. But unfortunately solving (5) is both numerically unstable and an NP-complete problem that requires an exhaustive enumeration of all $\binom{N}{K}$ possible combinations for the locations of the nonzero entries in \mathbf{s} .

Minimum ℓ_1 norm reconstruction: The compressive sensing surprise is that from $M \geq cK \log(N/K)$ iid Gaussian measurements we can exactly reconstruct K -sparse vectors and closely approximate compressible vectors stably with high probability via the ℓ_1 optimization [1, 2]

$$\hat{\mathbf{s}} = \operatorname{argmin} \|\mathbf{s}'\|_1 \quad \text{such that } \Theta \mathbf{s}' = \mathbf{y}. \quad (6)$$

This is a convex optimization problem that conveniently reduces to a linear program known as *basis pursuit* [1, 2] whose computational complexity is about $O(N^3)$.

To summarize, a compressive sensing data acquisition system consists of random measurements based on Φ followed by linear programming reconstruction to obtain \mathbf{x} .

6 Discussion

6.1 Geometrical interpretation

The geometry of the compressive sensing problem in \mathbb{R}^N helps us visualize why ℓ_1 reconstruction succeeds where ℓ_2 reconstruction fails. First, note that the set of all K -sparse vectors \mathbf{s} in \mathbb{R}^N is a highly nonlinear space consisting of all K -dimensional hyperplanes that are aligned with the coordinate axes (see Figure 2(a)). Thus, sparse vectors live close to the coordinate axes in \mathbb{R}^N .

To visualize why ℓ_2 reconstruction fails, note that the translated null space $\mathcal{H} = \mathcal{N}(\Theta) + \mathbf{s}$ is of dimension $(N - M)$ and is oriented at a random angle due to the randomness in the matrix Θ . See Figure 2(b), but beware that in practice $N, M, K \gg 3$, so you need to extrapolate your intuition to high dimensions. The ℓ_2 minimizer $\hat{\mathbf{s}}$ from (4) is the point on \mathcal{H} closest to the origin. We can find this point by blowing up a hypersphere (the ℓ_2 ball) until it touches \mathcal{H} . Due to the random orientation of \mathcal{H} , the closest point $\hat{\mathbf{s}}$ will with high probability live away from the coordinate axes and hence will be neither sparse nor close to the correct answer \mathbf{s} .

In sharp contrast to the ℓ_2 ball, the ℓ_1 ball in Figure 2(c) is “pointy” with the points aligned along the coordinate axes (and it becomes pointier as the ambient dimension N grows). Therefore, when we blow up the ℓ_1 ball, it will first touch the translated null space \mathcal{H} at a point near the coordinate axes, which is precisely where the sparse vector \mathbf{s} is located.

6.2 Analog signals

While we have focused on discrete-time signals \mathbf{x} , compressive sensing also applies to analog signals $x(t)$ that can be represented sparsely using just K out of N possible elements from some continuous basis or dictionary $\{\psi_i(t)\}_{i=1}^N$. While each $\psi_i(t)$ may have large bandwidth (and hence a high Nyquist rate), the signal $x(t)$ has only K degrees of freedom, and we can apply the above theory to measure it at a rate below Nyquist. An example of a practical analog compressive sensing system — a so-called “analog-to-information” converter — is given in [7]; there are interesting connections to [8].




Figure 2: (a) A sparse vector s lies on a K -dimensional hyperplane aligned with the coordinate axes in \mathbb{R}^N and thus close to the axes. (b) Compressive sensing recovery via ℓ_2 minimization does not find the correct sparse solution s on the translated nullspace (green hyperplane) but rather the non-sparse vector \hat{s} . (c) With enough measurements, recovery via ℓ_1 minimization does find the correct sparse solution s .

7 Practical Example

Consider the “single-pixel” compressive digital camera of Figure 3(a) that directly acquires M random linear measurements without first collecting the N pixel values [9]. The incident light-field corresponding to the desired image x is not focused onto a CCD or CMOS sampling array but rather reflected off a digital micromirror device (DMD) consisting of an array of N tiny mirrors. (DMDs are found inside many computer projectors and projection televisions.) The reflected light is then collected by a second lens and focused onto a single photodiode (the single pixel). Each mirror can be independently oriented either towards the photodiode (corresponding to a 1) or away from the photodiode (corresponding to a 0). Each measurement y_j is obtained as follows: The random number generator (RNG) sets the mirror orientations in a pseudorandom 0/1 pattern to create the measurement vector ϕ_j . The voltage at the photodiode then equals y_j , the inner product between ϕ_j and the desired image x . The process is repeated M times to obtain all of the entries in y .

Figure 3(b) and (c) illustrate a target object and an image \hat{x} taken by a prototype single-pixel camera [9] using about 60% fewer random measurements than reconstructed pixels. Here the reconstruction was performed via a total variation optimization [1], which is closely related to ℓ_1 reconstruction in the wavelet domain.

A major advantage of the single-pixel, compressive sensing approach is that this camera can be adapted to image at wavelengths where it is difficult or expensive to create a large array of sensors. It can also acquire data over time to enable video reconstruction [9].

8 Conclusions

In this lecture we have learned that sampling is not the only way to acquire signals. When the signals of interest are compressible or sparse, it can be more efficient and streamlined to employ random measurements and optimization in order to acquire only the measurements we need. We




Figure 3: (a) Single-pixel, compressive sensing camera. (b) Conventional digital camera image of a soccer ball. (c) 64×64 black-and-white image \hat{x} of the same ball ($N = 4096$ pixels) recovered from $M = 1600$ random measurements taken by the camera in (a). The images in (b) and (c) are not meant to be aligned.

have also learned that for reconstruction our old friend least squares fails us, and that we need to look to other flavors of convex optimization like linear programming (see also [4, 10]). Compressive sensing is still a nascent field, and many interesting research questions remain open. A fairly exhaustive set of references on the current research directions is available at dsp.rice.edu/cs.

References

- [1] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [2] D. Donoho, “Compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, April 2006.
- [3] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [4] J. Tropp and A. C. Gilbert, “Signal recovery from partial information via orthogonal matching pursuit,” April 2005, www-personal.umich.edu/~jtropp/papers/TG05-Signal-Recovery.pdf.
- [5] R. G. Baraniuk, M. Davenport, R. DeVore, and M. B. Wakin, “The Johnson-Lindenstrauss lemma meets compressed sensing,” 2006, dsp.rice.edu/cs/jlcs-v03.pdf.
- [6] D. Baron, M. B. Wakin, M. Duarte, S. Sarvotham, and R. G. Baraniuk, “Distributed compressed sensing,” 2005, dsp.rice.edu/cs/DCS112005.pdf.
- [7] S. Kirolos, J. Laska, M. Wakin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. G. Baraniuk, “Analog-to-information conversion via random demodulation,” in *IEEE Dallas Circuits and Systems Workshop*, Oct. 2006.
- [8] M. Vetterli, P. Marziliano, and T. Blu, “Sampling signals with finite rate of innovation,” *IEEE Trans. Signal Processing*, vol. 50, no. 6, pp. 1417–1428, June 2002.
- [9] D. Takhar, V. Bansal, M. Wakin, M. Duarte, D. Baron, J. Laska, K. F. Kelly, and R. G. Baraniuk, “A compressed sensing camera: New theory and an implementation using digital micromirrors,” in *Proc. Comput. Imaging IV at SPIE Electronic Imaging*, San Jose, Jan. 2006.

- [10] J. Haupt and R. Nowak, “Signal reconstruction from noisy random projections,” *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 4036–4048, Sept. 2006.

Acknowledgements: This work was supported by grants from NSF, DARPA, ONR, AFOSR, and the Texas Instruments Leadership University Program. Thanks to the Rice DSP group and Ron DeVore for many enlightening discussions. Thanks to Justin Romberg for help with the reconstruction in Figure 3.

Author: Richard G. Baraniuk is the Victor E. Cameron Professor of Electrical and Computer Engineering at Rice University. His research interests include multiscale analysis, inverse problems, distributed signal processing, and sensor networks. For his research he has received National Young Investigator awards from the US National Science Foundation and Office of Naval Research and achievement awards from Rice University and the University of Illinois. For his teaching he has received the C. Holmes MacDonald National Outstanding Teaching Award from Eta Kappa Nu and the George R. Brown Award for Superior Teaching at Rice University three times. He was elected a Fellow of the IEEE in 2001.