Tutorial on Compressive Sampling

Michael Lamoureux*
University of Calgary, Calgary, AB
mikel@math.ucalgary.ca

Summary

This is a brief tutorial on compressive sampling, a recently developed technology originating with mathematicians Candes (Caltech), Tao (UCLA), and others. This method greatly reduces the number of digital samples required to reconstruct certain important physical signals. The key ideas of this technology are presented here, along with some indications of possible applications to geophysical problems.

Introduction

Compressive sampling is a recent development in digital signal processing that offers the potential of high resolution capture of physical signals from relatively few measurements, typically well below the number expected from the requirements of the Shannon/Nyquist sampling theorem. This technique combines two key ideas: sparse representation through an informed choice of linear basis for the class of signals under study; and incoherent (eg. pseudorandom) measurements of the signal to extract the maximum amount of information from the signal using a minimum amount of measurements. Mathematical techniques required to implement compressive sampling include the development of novel types of linear bases (eg. wavelet, curvelet, etc.), L1 optimization to recover sparse representations, and design of optimal dual measurements.

Theoretical Development

It is well understood that many recorded signals such as music, photos, medical images, seismic data, can be stored in compressed form by expressing them in terms of a linear basis, wherein many of the coefficients in the representation are zero or small enough to be ignored. This is the technology behind wavelet compression, JPEGs, and MP3 players, among others. Typical examples of such compressive-type bases include 2D wavelets for images, localized sinusoids for music, fractal-type waveforms for spiky reflectivity data, and curvelets for wavefield propagation.

An optimal sparse representation in a given basis is obtained by performing a constrained L1 optimization over the linear coefficients that appear in the representation of the signal. That is, given some signal f and basis elements $\psi_1, \psi_2, ..., \psi_N$, a minimization is performed, as

$$\min \sum_{j=1}^{N} |a_j| \text{ subject to } ||f - \sum_{j=1}^{N} a_j \psi_j|| < \epsilon.$$

The L1 minimization tends to concentrate the energy of the signal on to a few non-zero coefficients a_j, unlike the least squares (L2 minimization) which tends to spread the energy around. This results

in the sparse representation. By replacing the absolute value of the a_j with the difference of the positive and negative parts, the L1 function becomes a linear objective which is tractable even in large dimensions, by using the simplex or interior point methods of linear programming.

Compressive sampling takes this idea one step further by creating a measurement system for physical systems so that the real signal itself can be recorded in compressed form "on the fly." This bypasses the necessity to first capture digitally the full signal at high resolution and high data rate, and then compress "in the box."

A key step is the creation of measurement vectors $\phi_1, \phi_2, ..., \phi_K$ for taking physical measurements on the signal in the form of inner products of the signal with the measurement vectors, $y_k = \langle f,\phi_k \rangle$. The measurement vectors are carefully designed to extract the maximum amount of information from a generically sparse vector in the given basis system. The optimization problem is replaced by a linearly constrained problem where the measurements of the signal must match the measurements on the representative solution. That is, one solves

$$\min \sum_{j=1}^{N} |a_j| \text{ subject to } y_k = \langle \sum_{j=1}^{N} a_j \psi_j, \phi_k \rangle, k = 1 \ldots K$$

One of the deep results in the theory is that for a sparse signal of order S, only on the order of K = S log(N) measurement vectors are needed. In a successful compressive system, one designs for K << N.

The measurement vectors, when organized into a KxN, matrix, must satisfy the restricted isometry condition: that is, for all T ≤ S, all KxT submatrices Φ and all T-vectors x, one has

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

That is, each submatrix is approximately an isometry. The Uniform Uncertainty Principle (UUP) of Candes and Tao tells us that for δ sufficiently small, the constrained L1 reconstruction is almost always exact on sparse signals, with high probability.

Finding such measurement vectors is a challenge and forms an active area of research. One difficulty is that for practical problems, the dimensions can be very high. Some successful approaches include: randomly choose K rows from an NxN orthogonal matrix and normalize the columns; randomly choose N unit vectors in K-dimensional space, and organize into the KxN matrix; form a matrix with randomly chosen Gaussian entries. With high probability, these choices work.

Examples

As an illustrative, but elementary example, consider the case of a one-dimensional signal that is sparse in the frequency domain. We assume a function expressible in the form of a sum of a small number of sinusoids,

$$f(t) = a_{j_1} \sin(\omega_{j_1} t) + a_{j_2} \sin(\omega_{j_2} t) + \cdots + a_{j_S} \sin(\omega_{j_S} t),$$

where the coefficients a_j are the amplitudes and the w_j are the frequencies of the sinusoids. From the Shannon sampling theorem, the number of regular samples of the signal required for perfect reconstruction will depend on the bandwidth of the signal, and the frequency resolution desired.

The sparse basis for this type of signal would be collections of sinusoids of the form $\sin(w_j t)$, j=1...N, where the frequencies span the bandwidth at the desired resolution. A suitable incoherent measurement system for this basis is to select random samples in the time domain, obtaining

measurements $y_k = f(t_k)$, where the t_k, k=1...K, are selected randomly. The L1 optimization problem is the constrained minimization over the variables a_j, expressed in the form

$$\min \sum_{j=1}^{N} |a_j| \text{ subject to } y_k = \sum_{j=1}^{N} a_j \sin(\omega_j t_k), k = 1 \ldots K.$$

Splitting the a_j into positive and negative parts, this gives a linear programming problem with 2N variables and K linear constraints.

For a numerical example, we take the signal f(t) with S=2 sinusoids in it, with frequency content in the band 0 to 10Hz, and seek a resolution of 0.1Hz. At this bandwidth and resolution, some 100 sinusoids are required in the basis. Compressive sampling requires on the order of K = S log N random samples, or about K = 2 log 100 = 9.2 samples. Figure 1 shows a typical waveform, with 9 random samples, and the recovered coefficients a_j. Notice that in this example, the reconstructed signal is an exact replica of the original, and thus the frequency content, as expressed in the coefficients a_j, is also recovered exactly.

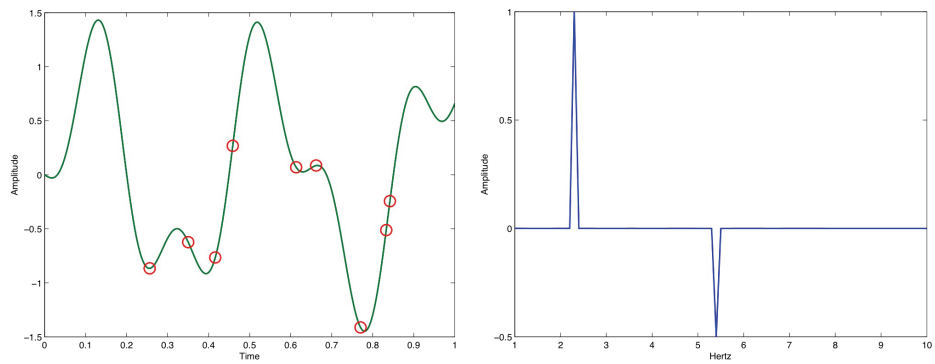

Figure1: On the left, random samples of the function f(t) = sin(2.3•2πt) - .5sin(5.4•2πt).
On the right, the coefficients of the best L1 fit: a perfect reconstruction.

As a comparison, Figure 2 shows the analogous least square reconstruction of the signal and corresponding frequency content in terms of the recovered coefficients a_j. In the least squares case, the reconstructed signal is far from exact, although it is an exact match at the sample points. Also, the frequency content of the reconstruction is significantly different from the original. The reason for this failure is that the least squares approach would require many more than 9 samples for an accurate reconstruction. For instance, the Shannon sampling theorem would predict on the order of 200 uniform samples to recover the frequency content to 10Hz bandwidth at 0.1 Hz resolution.

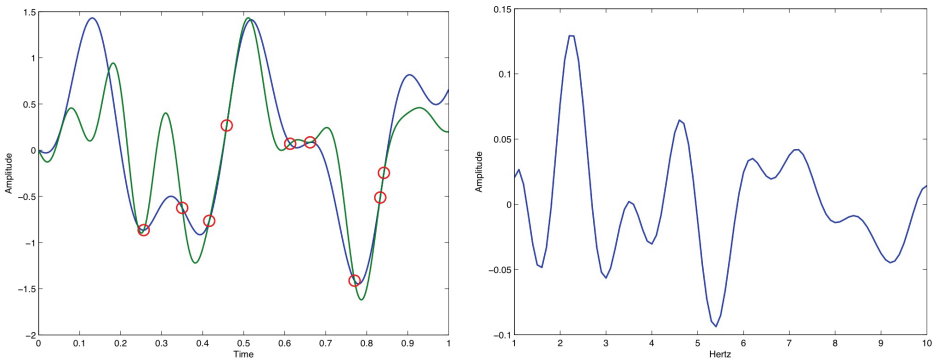

Figure 2: On the left, random samples of f(x) = sin(2.3t) - .5sin(5.4t) and the least squares fit. On the right, the corresponding frequency coefficients: not a perfect fit.

In Figure 3, we borrow an example from Candes of the compressive sampling and reconstruction of a recorded angiogram. In this case, the measurements are Fourier transforms of the radial slices through the image, analogous to slices of a CAT scan. A weighted L1 optimization is used to recover the image from these measurements. As a comparison, a back projection reconstruction is shown side-by-side with the L1 reconstruction: notice the higher resolution of the L1 construction, and the great reduction of artifacts as compared to the back projection.

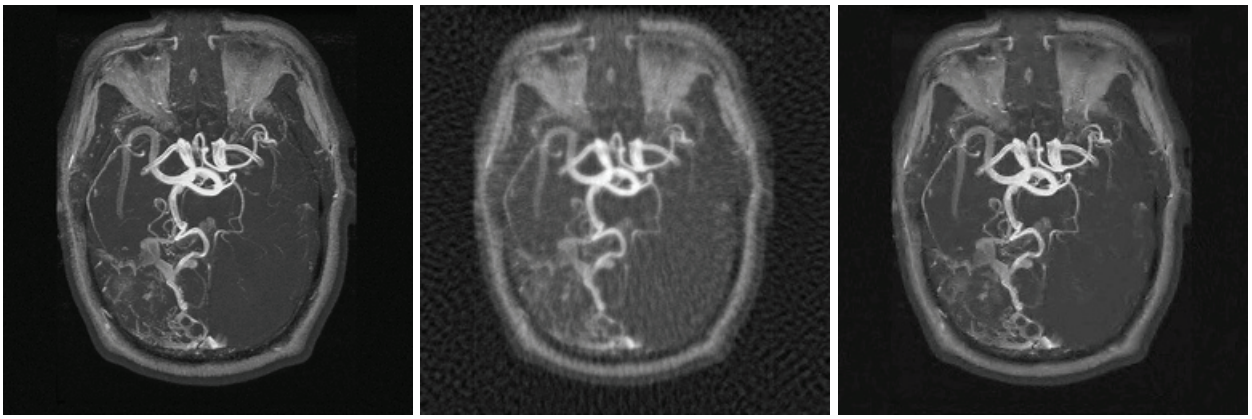

Figure 3: Reconstruction of an angiogram.
Left: original image. Middle: back projection. Right: L1 optimal reconstruction.

Possible Geophysical Applications

It is known that seismic wavefields have sparse representations in the curvelet domain; other seismic data can also be represented sparsely in suitable domains (Herrmann 2005). In a typical seismic survey, the data is vastly undersampled in the spatial domain and yet accurate images can be recovered from these measurements. Compressive sampling and the UUP may give a theoretical foundation for understanding this phenomena and lead to better algorithms for processing the seismic data.

Conclusions

Compressive sampling offers a method of recording high resolution physical data at vastly undersampled rates, using techniques of sparse basis representation, L1 optimization, and carefully designed measurement systems based on the Uniform Uncertainty Principle. With previous success in medical imaging technologies, we expect it to have uses in seismic processing as well.

Acknowledgements

This research is supported by an NSERC Discovery Grant, MITACS Inc., and the sponsors of the POTSI and CREWES projects.

References

Candes, E., 2006, Compressive sampling, Proceedings of the International Congress of Mathematicians: Madrid.

Candes, E. and Romberg, J., 2006, Sparsity and incoherence in compressive sampling: Inverse Problems, **23**, 969–985.

Candes, E., Romberg, J., and Tao, T., 2006, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information: IEEE Trans. Inform. Theory, **52**, 489–509.

Candes, E., Wakin, M., and Boyd, S., 2007, Enhancing sparsity by reweighted L1 minimination: Technical Report, Caltech.

Candes, E. and Wakin, M., 2007, People hearing without listening: an introduction to compressive sampling: IEEE Signal Processing Magazine, to appear.

Herrmann, F., 2005, Seismic deconvolution by atomic decomposition: a parametric approach with sparseness constraints: Integrated Computer-Aided Engineering, **12**:1, 69–90.

Nocedal, J. and Wright, S., 2006, Numerical Optimization, Springer Series in Operations Research, Springer.