Exact reconstructions from surprisingly little data

Rick Chartrand, rickc@lanl.gov

Recent results [1, 2] have shown that it is possible to obtain exact reconstruction of a large class of signals from far less data than is suggested by traditional sampling theory. Now we have shown how to extend the method to decrease the necessary amount of data still further. The underlying phenomenon is very general, and applies to signals of all types. Decoding highly corrupted bit streams, image or 3-D object reconstruction from very few radiographs, and network tomography are just a few of the many possible applications in all aspects of LANL's mission.

An old result in signal processing describes a situation where one can reconstruct a continuously varying signal exactly from only finitely many samples. The Shannon sampling theorem says that if the signal has a largest frequency component, then it can be reconstructed exactly if it is sampled at least as often as twice that frequency (known as the Nyquist rate). For many discrete signals, however, this would require one to sample the entire signal. Generally, the Nyquist rate should be regarded as an upper bound of how much data is necessary to reconstruct a signal; often, much less data is required.

In the new theory, known as "compressed sensing" or "compressive sampling," the notion of a Nyquist frequency is replaced with the notion of *sparsity*. The simplest example of a sparse signal is one that is zero in all but a few places. A more general example is an image that is well approximated by a wavelet expansion with not too many nonzero wavelet coefficients. This includes any image that is compressible, and thus almost any image of the real world.

For an example from radiography, we consider the Shepp-Logan phantom, a synthetic test image shown in the first figure. It is sparse in the sense of having piecewise-constant intensity, so that the gradient of its intensity is zero except at the boundaries of the objects in the image. In the context of CT X-ray inversions, typically a radiograph would be taken at 180 or 360 different angles in order to expect to get a good reconstruction of the object. The sparsity of this phantom makes it possible to reconstruct it *exactly* from only 10 radiographs.

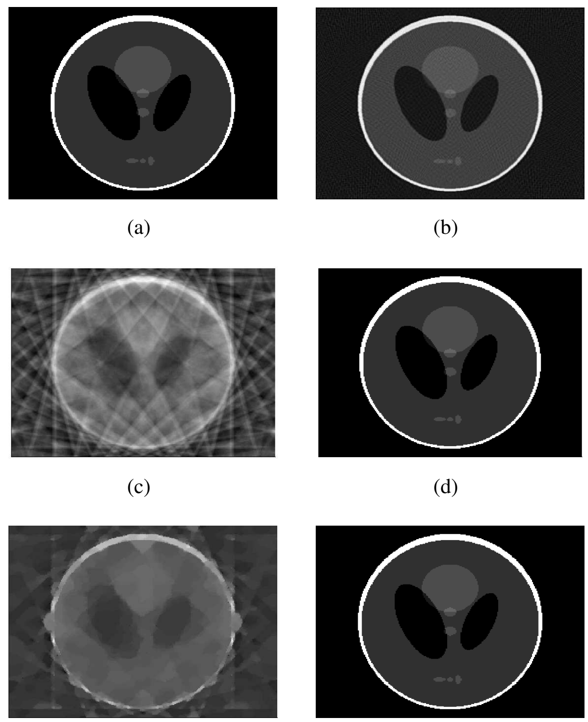

(a) (b) (c) (d) (e) (f)

(a) The Shepp-Logan phantom. (b) Traditional (backprojection) reconstruction from radiographic projections along 180 angles is good but not perfect. (c) Traditional reconstruction from 18 projections is terrible. (d) ℓ^1 minimization gives an exact reconstruction with 18 projections. (e) ℓ^1 minimization gives a poor reconstruction from 10 projections. (f) ℓ^p minimization with $p = 1/2$ gives an exact reconstruction with 10 projections.

We explain how this is done in a general setting. Suppose f is a signal that is sparse in some sense. We can regard f as a vector with n components; if f is not a one-dimensional signal (like an image or network), we just string out the components of f into a vector. We formulate sparsity "in some sense" as there being an $n \times n$ matrix V such that Vf is nonzero in not too many places, say k of them. (V could be a change-of-basis matrix, or a discrete approximation of a derivative operator.) Let M be a measurement matrix of size $m \times n$, so that Mf would be a vector containing m measurements of the signal f. (What sort of measurements will be described shortly.) We then seek our reconstruction as the solution of the following optimization problem:

$$\min_u \|Vu\|_p^p, \text{ subject to } MVu = MVf. \qquad (1)$$

Here we use the ℓ^p norm, defined by $\|x\|_p^p = \sum_{i=1}^n |x_i|^p$, and $0 < p \leq 1$. In the compressed sensing literature, $p = 1$. In this case, it is proven [1] that the solution to (1) will be exactly $u = f$, provided the number of measurements satisfies $m \geq Ck \log n$, for some not-large constant C. The key is that the required number of measurements nearly depends only on the sparsity k, and can be a small fraction of the signal size n.

There are many instances where surprisingly few measurements are required to get exact reconstructions when $p = 1$. Our contribution [3, 4] is to show that using $p < 1$ allows even fewer measurements, as little as half as many as when $p = 1$.

Returning to the phantom example, the measurements we use are samples of the Fourier transform, taken along radial lines through the origin in frequency space. By the Fourier-slice theorem, this is equivalent to sampling radiographic projections onto detectors perpendicular to these lines. For $p = 1$, 18 projections are enough for an exact reconstruction, in the sense of every pixel value being with 10^{-13} of the correct value. If we use $p = 1/2$, 10 suffice. By comparison, a traditional backprojection reconstruction gives only an approximate reconstruction with 180 or even 360 projections, and a terrible one with 18. Similarly, the $p = 1$ reconstruction with 10 views is poor.

Not all measurement matrices will exhibit this phenomenon. What is required is an *incoherence* property between the measurement basis and the basis in which the signal is sparse. The result will be a generalization of Heisenberg's uncertainty principle, in which a signal cannot be sparse in both bases simultaneously. Then the sparsity of the signal ensures that the measurements give as much information as possible.

A universal way to (nearly) guarantee that the measurements will have the incoherence property is to make them random. For example, the elements of the measurement matrix M can be chosen from a Gaussian distribution, or to have value ± 1 with equal probability. Then, regardless of the sparse representation V, M will have the necessary incoherence property with overwhelming probability (tending to 1 as $n \to \infty$). This is highly counterintuitive: such random measurements will appear to turn the signal into gibberish, yet will maximize the likelihood of obtaining an exact reconstruction.

We are exploring applications to 3-D radiography, correcting errors in encypted signals [4], and network tomography. These and other applications considered in the literature only begin to harness the potential of these methods.

Acknowledgements

Los Alamos Report LA-UR-06-6184. Funded by the Department of Energy at Los Alamos National Laboratory under contract DE-AC52-06NA25396.

References

[1] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*, 52, 2006.

[2] D. L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52:1289–1306, 2006.

[3] R. Chartrand. Exact reconstruction of sparse signals via nonconvex minimization. Submitted.

[4] R. Chartrand. Nonconvex compresseed sensing and error correction. Submitted.