Compressed Sensing

David Donoho, Statistics Department, Stanford University

Suppose x is an unknown vector in R^m (depending on context, a digital image or signal); we plan to acquire data and then reconstruct. Nominally this 'should' require m samples. But suppose we know *a priori* that x is compressible by transform coding with a known transform, and we are allowed to acquire data about x by measuring n general linear functionals – rather than the usual pixels. If the collection of linear functionals has certain properties, and we allow for a degree of reconstruction error, the size of n can be dramatically smaller than the size m usually considered necessary. Thus, certain natural classes of images with m pixels need only $n = O(m^{1/4} \log(m))$ nonadaptive nonpixel samples for faithful recovery, as opposed to the usual m pixel samples.

Underlying our results is a theoretical framework based on the theory of optimal recovery, the theory of n-widths, and information-based complexity. The basic results concern properties of ℓ^p balls in high-dimensional Euclidean space: the Gel'fand n-widths of such balls, new families of near-optimal subspaces for Gel'fand n-widths, and new algorithms for processing information derived from near-optimal subspaces.