Towards Deterministic Compressed Sensing

Jeffrey D. Blanchard, Department of Mathematics and Statistics, Grinnell College, Grinnell, IA

Over the past decade, compressed sensing has delivered significant advances in the theory and application of measuring and compressing data. Consider capturing a 10 mega pixel image with a digital camera. Emailing an image of this size requires an unnecessary amount of storage space and bandwidth. Instead, users employ a standard digital compression scheme, such as JPEG, to represent the image as a 64KB file. The compressed image is completely recognizable even though the dimension of the compressed version is a tiny fraction of the original 10 million dimensions. Compressed sensing takes this mathematical phenomenon one step further. Is it possible to capture the pertinent information, such as the 64KB image, without first measuring the full 10 million pixel values? If so, how should we perform the measurements? If we capture the important information, can we still reconstruct the image from this limited number of observations? Compressed sensing exploded in 2004 when Donoho (1,2) and Candes and Tao (3) definitively answered these questions by incorporating randomness in the measurement process. Since engineering a truly random process is impossible, a major open problem in compressed sensing is the search for deterministic methods for sparse signal measurement that capture the relevant information in the signal and permit accurate reconstruction. In this issue of PNAS, Monajemi et al. (4) provide a major step forward in understanding the potential for deterministic measurement matrices in compressed sensing.

Capturing digital images on a camera is simple; however, there are many applications where the measurement process has a much greater underlying cost. Magnetic Resonance Imaging (MRI) is a prime example of a high-impact compressed sensing application. For most MRI examinations, a patient is required to lie still in a confined space for roughly 45 minutes. In some situations, compressed sensing has generated diagnostic-quality MR images using only 10% as many measurements (5). MRI is only a single example of compressed sensing applications which extend well beyond imaging and include computed tomography, electro-cardiography, multispectral imaging, seismology, analog to digital conversion, radar, X-ray holography, astronomy, DNA sequencing, genotyping, and more (6).

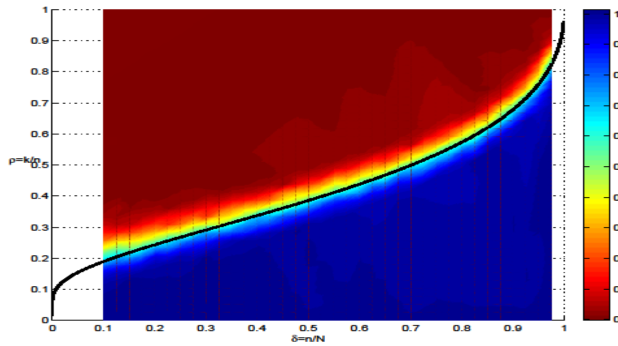

Fig. 1 Universality Hypothesis: Random Fourier Measurements (9, by permission of the Royal Society). Let k be the number of nonzero entries in a signal of length N, and n the number of linear measurements observed through the $n \times N$ measurement matrix with $k < n < N$. Two ratios, $\rho = k/n$ (vertical axis) and $\delta = n/N$ (horizontal axis) define the compressed sensing phase space in the unit square $0 < \rho, \delta < 1$. The black curve is the Gaussian measurement matrix phase transition defined by a function $\rho^*(\delta)$: if $\rho < \rho^*(\delta)$ then l_1 minimization successfully reconstructs almost every signal yet almost always fails when $\rho > \rho^*(\delta)$. Shaded attribute: fraction of realizations in which l_1 minimization successfully reconstructs a signal measured by a random subset of n rows of a Fourier matrix.

Traditional signal processing procedures measure the full signal directly and apply standard compression routines for storage or transmission. When needed, the original signal can be reconstructed by inverting the linear compression procedure. Compressed sensing transfers the workload from the measurement process to the signal reconstruction. While the measurement process remains linear, the reduced number of measurements forces a highly nonlinear reconstruction process.

Rather than taking point measurements of the entire signal, compressed sensing uses more sophisticated measurement schemes which acquire information throughout the signal and mix the information into relatively few numerical values. Decoding these complicated measurements from the underdetermined system of equations is therefore considerably more challenging than most other signal reconstruction techniques. In fact, because the system is underdetermined and at least one signal could have generated the linear measurements, there exist infinitely many signals that generate the exact same measurements. Compressed sensing relies on the assumption that the original signal has a low information content compared to its physical dimension. Typically, the low information content is interpreted as sparsity where a signal is "sparse" when the number of nonzero entries in the signal's digital representation is dramatically smaller than its ambient dimension. This assumption is justified by the vast literature on signal processing where, for example, images are known to be sparse in, say, the wavelet domain. That is, images have very few large coefficients when represented in a wavelet basis and can be accurately approximated using only these large coefficients.

A traditional technique for selecting a particular solution of an underdetermined linear system is to form the least squares solution by simply multiplying the measurement values by the pseudo-inverse of the measurement matrix. Solving the least squares problem produces the signal with the minimum l_2 norm (the square root of the sum of the squares of the entries) among the infinite set of signals which produce the same measurements. Minimizing the l_2 norm returns a signal whose total energy is distributed throughout the signal and therefore fails to meet the sparsity assumption. In the compressed sensing regime, one wishes to obtain the signal with the fewest number of nonzero entries. This presents a combinatorial optimization problem whose naïve solution is as hard as the famous traveling salesman problem. For some time it has been observed that replacing the least squares problem with minimizing the l_1 norm (the sum of the absolute values of the entries) produces a sparse solution. The frenzy surrounding compressed sensing began when Donoho (1, 2, 7), Candes, Romberg, and Tao (3, 8) proved sufficient conditions which ensure the solution to the l_1 minimization problem coincides with the sparsest solution from the combinatorial optimization problem. The beauty of this finding is that the intractable combinatorial optimization is replaced by a tractable convex optimization problem.

Both the geometric condition of Donoho, namely neighborliness of an associated polytope, and the linear algebraic condition of Candes, Romberg, and Tao, known as the restricted isometry property, are deterministic conditions. However, checking the validity of either of these conditions is also a combinatorial problem. This theoretical obstacle was overcome by analyzing random matrices such as Gaussian matrices whose entries are drawn independently and identically from a normal distribution. They proved that measurement matrices drawn from certain random matrix ensembles captured the pertinent information with very few measurements. Moreover, sparse signals could then be accurately reconstructed by solving the tractable, convex l_1 minimization problem. This combined the measurement and compression processes into a single random measurement process.

In this issue, Monajemi et al. (4) bring together two major lines of research in compressed sensing, the Universality Hypothesis and the search for deterministic measurement matrices. The Universality Hypothesis, formally stated by Donoho and Tanner (9), claims many families of random matrices exhibit the same signal recovery performance via l_1 minimization as the Gaussian matrices. Donoho and Tanner (1, 2, 10) formally proved the Gaussian measurement ensemble exhibits a phase transition where the probability of successful recovery abruptly changes from 1 to 0 as a function of two parameters defining the compressed sensing problem (see Fig. 1). The Universality Hypothesis was empirically established by Donoho and Tanner (9) for several random matrix ensembles in that their observed success and failure of signal reconstruction under l_1 minimization matches the performance of the Gaussian ensemble. A proof of the Universality Hypothesis would extend nearly all of the compressed sensing theory established for Gaussian ensembles to a much larger class of random matrix ensembles. In mid 2012, Bayati, Lelarge, and Montanari announced a major advance proving the Universality Hypothesis for a wide class of random matrices.

An independent line of research seeks to remove randomness from the compressed sensing regime. The main tool in this pursuit is the restricted isometry property (RIP) of Candes and Tao (3). DeVore presented explicit constructions of matrices satisfying the RIP, but these constructions require the number of measurements to scale with the square of the sparsity (11). In contrast, the Gaussian matrices require the measurements to scale linearly with the sparsity plus a minor logarithmic penalty. Several other groups have continued this line of work improving on the requisite number of measurements, for example (12, 13); so far no theoretical results for deterministic matrices match the optimal measurement rate achieved by Gaussian matrices. Alternative approaches for analyzing deterministic matrices (14-16) established theoretical justification for their successful sparse signal recovery. This important line of research has not yet established deterministic compressed sensing performance on par with the random measurement models.

In the companion article (4), Monajemi et al. present overwhelming empirical evidence that the Universality Hypothesis should now also include many deterministic measurement ensembles. Through considerable experimental analysis, they have shown that nine deterministic matrix families exhibit the same phase transition behavior as Gaussian matrices when the signal is reconstructed via l_1-minimization. Most of these deterministic matrices come from the work discussed in the preceding paragraph. The potential advantages of deterministic measurements in applications cannot be understated. Not only does randomness present considerable engineering challenges for implementation, but dense random matrices consume large amounts of memory and require computationally expensive matrix multiplications. Most of the reported deterministic matrix families avoid the need to store the entire matrix and boast fast matrix multiplications akin to the Fast Fourier Transform, providing massive computational gains. The incorporation of deterministic measurement matrices in the Universality Hypothesis not only bolsters the importance of this conjecture, it opens a wide range of new possibilities and intriguing questions for compressed sensing.

1. D. L. Donoho (2006) For most large underdetermined systems of equations, the minimal l1-norm solution is also the sparsest solution. Comm. Pure Appl. Math., 59(6):797-829.
2. D. L. Donoho (2006) For most large underdetermined systems of equations, the minimal l1-norm near-solution approximates the sparsest near-solution. Comm. Pure Appl. Math., 59(7):907-934.
3. E. J. Candes ,T. Tao (2005) Decoding by linear programming. IEEE Trans. Inform. Theory, 51(12):4203-4215.
4. H. Monajemi, S. Jafarpour, M. Gavish, Stat330/CME362 Collaboration, D. L. Donoho (2012) Deterministic matrices matching the compressed sensing phase transitions of Gaussian random matrices. Proc. Natl. Acad. Sci. USA, 1000(1000):1000-1001.
5. M. Lustig, D. L. Donoho, J. M. Santos, J. M. Pauly (2008) Compressed sensing MRI. IEEE Signal Processing Magazine, 25(2):72-82.
6. Rice Compressed Sensing Resources, http://dsp.rice.edu/cs.
7. D. L. Donoho (2005) Neighborly polytopes and sparse solution of underdetermined linear equations. Technical Report, Department of Statistics, Stanford University.
8. E. J. Candes, J. Romberg, T. Tao (2006) Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Trans. Inform. Theory, 52(2):489-509.
9. D. L. Donoho, J. Tanner (2009) Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. Phil. Trans. Royal Society A, 367(1906):4273-4293.
10. D. L. Donoho, J. Tanner (2005) Sparse nonnegative solution of underdetermined linear equations by linear programming. Proc. Natl. Acad. Sci. USA, 102(27):9446-9451 (electronic).
11. R. A. DeVore (2007) Deterministic constructions of compressed sensing matrices. J. Complexity, 23(4-6):918-925.
12. J. Bourgain, S. Dilworth, K. Ford, S. Konyagin, D. Kutzarova (2011) Explicit constructions of RIP matrices and related problems. Duke Math. J., 159(1):145-185.
13. H. Rauhut, J. Romberg, J. A. Tropp (2012) Restricted isometries for partial random circulant matrices. Appl. Comput. Harmon. Anal., 32(2):242-254.
14. J. Tropp (2008) On the conditioning of random subdictionaries. Appl. Comp. Harm. Anal., 25(1):1-24.
15. E. J. Candes , Y. Plan (2009) Near-ideal model selection by l1 minimization. Ann. Statist., 37(5A):2145-2177.
16. R. Calderbank, S. Howard, S. Jafarpour (2010) Construction of a large class of deterministic sensing matrices that satisfy a statistical isometry property. IEEE Selected Topics in Sig. Proc., 4(2):358-374.