Titles and Abstracts

Ron DeVore

Title: How should we measure performance in compressed sensing?

Abstract: We shall discuss possible ways to measure the performance of a compressed sensing system (encoder and decoder). In particular we shall introduce the concept of instance optimality. We shall then deduce that the most natural measure is instance optimality in probability. We then discuss which encoders and decoders have optimal performance with respect to this property.

Rich Baraniuk and Volkan Cevher

Title: Model-based compressive sensing

Compressive sensing (CS) is an alternative to Shannon/Nyquist sampling for acquisition of sparse or compressible signals that can be well approximated by just K<< N elements from an N-dimensional basis. Instead of taking periodic samples, we measure inner products with M<N random vectors and then recover the signal via a sparsity-seeking optimization or greedy algorithm. The standard CS theory dictates that robust signal recovery is possible from M=O(K log(N/K)) measurements. The implications are promising for many applications and enable the design of new kinds of analog-to-digital converters, cameras and imaging systems, and sensor networks.

While this represents significant progress from Nyquist-rate sampling, in this talk, we will demonstrate that it is possible to do even better by more fully leveraging concepts from state-of-the-art signal compression and processing algorithms. In many such algorithms, the key ingredient is a more realistic signal model that goes beyond simple sparsity by codifying the inter-dependency structure among the signal coefficients. We will present a new model-based CS theory that parallels the conventional theory and provides concrete guidelines on how to create model-based recovery algorithms with provable performance guarantees. To take practical advantage of the new theory, we integrate two relevant signal models -- wavelet trees and block sparsity -- into two state-of-the-art CS recovery algorithms and prove that they offer robust recovery from just M=O(K) measurements.

Arkadi Nemirovski

Title: L1-recovery: Verifiable Quality Guarantees and Efficient First Order Algorithms

In the talk, we outline our research plans and recent results related to:

(1) building efficiently computable upper and lower bounds on the largest s for which a given sensing matrix A is s-good, i.e., L1-recovery of every signal x with s nonzero entries from noiseless observations of Ax is exact;
(2) deriving error bounds for L1 recovery in the ``nonideal case'' (noisy observations, nearly s-sparse signal, imperfect L1 minimization) for s-good sensing matrices;
(3) ``computationally cheap'' deterministic and randomized first order algorithms for L1-recovery.

Joel Tropp

Title: Beyond Nyquist: Efficient sampling of sparse, bandlimited signals

Wideband analog signals routinely push state-of-the-art analog-to- digital conversion systems to their performance limits. In many applications, however, sampling at the Nyquist rate is inefficient because the signals of interest contain only a small number of frequencies relative to the bandwidth. For these type of sparse signals, other sampling strategies are possible.

This talk describes a new type of data acquisition system, called a random demodulator, that is constructed from robust, readily available components. Let K denote the total number of significant frequencies, and let W denote the bandwidth in Hz. New theoretical work, supported by empirical studies, establishes that the random demodulator requires just O(K*polylog(W)) samples per second to reconstruct any such signal. This sampling rate is exponentially lower than the Nyquist rate of W Hz. In contrast with Nyquist sampling, one must use nonlinear methods, such as convex programming, to recover the signal from the samples taken by the random demodulator. Rigorous algorithmic guarantees are also available.

Joint with R. Baraniuk, M. Duarte, J. Laska, and J. Romberg

Robert Calderbank

Title: Fast Reconstruction Algorithms for Deterministic Sensing Matrices and Various Applications

Abstract: Compressed Sensing aims to capture attributes of a sparse signal using very few measurements. The Restricted Isometry Property is the condition that the sensing matrix acts as as near isometry on all k-sparse signals. Candes and Tao showed that this condition is sufficient for sparse reconstruction and that random matrices, where the entries are generated by an iid Gaussian or Bernoulli process, satisfy the RIP with high probability. This approach treats all k-sparse signals equally likely, in contrast to mainstream signal processing where the filtering is deterministic, and the signal is described probabilistically. In the mainstream framework the sensing matrix is deterministic and it is required to act as a near-isometry on k-sparse vectors with high probability. We provide weak conditions that are sufficient to show that a deterministic sensing matrix satisfies this Statistical Restricted Isometry Property (STRIP). We also present applications to A/D conversion and wireless communication.

Justin Romberg & Karim Sabra
Georgia Institute of Technology, Atlanta

William A. Kuperman
Scripps Institution of Oceanography, San Diego

Title: Synthetic Aperture Compressive Sensing for Extracting Structured Scatterers in Complex Media: An MCM Application

Recent at-sea experiments have demonstrated that SONAR systems operating at lower frequency sonar could be used to excite the resonant acoustic signatures of proud and, more importantly, buried man-made targets. For this reason, the US Navy is currently investigating the application of *multi-bistatic* SONAR systems (i.e. where the source and receiver roles are performed by two different transducer arrays) using autonomous off-board platforms to create a wide, combined physical and synthetic aperture for mapping the *spatial and temporal characteristics* of the scattered field produced by the targets in the low to mid-frequency regime (<50kHz). Ideally, *multi-bistatic* SONAR systems would provide the additional spatial coverage and target "view points" to enhance target's detection but these systems would potentially generate very large amounts of data. The long-term goal of this project is to develop a compressive sensing methodology to minimize the number of multistatic acoustic measurements required to provide concurrent detection, classification and localization (DCL) of proud and buried targets in shallow ocean waveguides. A compressive sensing architecture is crucial for practical Navy applications with distributed sensor networks where each node (e.g. AUV) has a limited data-storage capacity and also due to low bit-rate currently available for robust underwater telecommunications. We will also discuss how recent results that link together compressed sensing and multiple channel estimation can be applied to *multi-bistatic* SONAR systems operating in complex environment such as shallow water waveguides.

Stan Osher

Title: Bregmanized methods for sparse reconstruction and restoration

Many of the issues in compressive sensing and image restoration reduce to variational problems, constrained or unconstrained, involving L1 type minimization. These include L1, TV, B_1,1, nonlocal TV etc. Bregman iteration, in various incarnations, seems to be appropriate for most of these. I'll discuss its unreasonable effectiveness and give results, theorems, and intuition as to why this is so.

Jerome Darbon

Title: Simple Compressive Algorithms for Parallel Many-core Architectures

Abstract: We consider the recovery of signal via compressive sensing where the signal itself or its gradient are assumed to be sparse. This amounts to solve a l^1 or a Total Variation minimization problem.

We propose minimization algorithms specifically designed to take advantage of shared memory, vectorized, parallel and many-core microprocessors such as the Cell processor, new generation Graphics Processing Units (GPUs) and standard vectorized multi-core processors (e.g. standard quad core CPUs). Besides their implementations are easy. We also give evidence of the efficiency of our approach and compare the algorithm on the three platforms, thus exhibiting pros and cons for each of them.

Larry Carin

Title: On the relationship between compressive sensing and random sensor arrays

Abstract: Random sensor arrays are examined from a compressive sensing (CS) perspective. It is demonstrated that the natural random-array projections manifested by the media Green's function are consistent with the projection-type measurements associated with CS. This linkage allows the use of existing CS theory to quantify the performance of random arrays, of interest for array design.

The analysis demonstrates that the CS theory is applicable to arrays in vacuum as well as in the presence of a surrounding media; further, the presence of a surrounding media with known properties may be used to improve array performance.

Donald Goldfarb

Part I

Title: Fixed point and Bregman iterative methods for matrix rank minimization

Abstract: The linearly constrained matrix rank minimization problem is the matrix analog of the compressive sensing recovery problem. The linearly constrained nuclear norm minimization (NNM)problem is a convex relaxation of this problem, which can be cast as a semidefinite programming (SDP) problem. Unfortunately, these SDPs are expensive to solve when the matrices are large. In this talk, we present and analyze fixed point and Bregman iterative algorithms for solving the NNM problem.
By using a homotopy approach together with an approximate singular value decomposition procedure, we get a very fast, robust and powerful algorithm that can solve very large matrix rank minimization problems. Our numerical results on randomly generated and real matrix completion problems demonstrate that this algorithm is much faster and provides much better recoverability than SDP solvers such as SDPT3.

Part II (with Zaiwen Wen)

Title: A fast algorithm for sparse reconstruction based on shrinkage, subspace optimization and continuation

Abstract. We describe a fast algorithm for sparse reconstruction. The algorithm is divided into two stages that are performed repeatedly. In the first stage, "shrinkage" yields an estimate of the subset of variables likely to be nonzero in an optimal solution. Restricting the decision variables to this subset and fixing their signs at their current values results in a smooth quadratic problem that is solved in the second phase.
Our method also embeds this basic two-stage algorithm in a continuation (homotopy) approach. Our implementation of this method exhibits state-of-the-art performance both in terms of its speed and its ability to recover sparse signals. It can even recover signals that are not as sparse as required by current compressive sensing theory.

Wotao Yin

Title: Enhanced Compressed Sensing based on Iterative Support Detection

Abstract: We demonstrate that the recovery rate of basis pursuit on fast decaying signals can be enhanced by applying a novel iterative support detection strategy. Preliminary theoretical and experimental results, as well as the limitation of the strategy, are presented. In addition, we spend fifteen minutes briefing the CS-related research and results of the CAAM Department at Rice University.

Justin Romberg

Title: Multiple channel estimation and Dynamic Updating

In this talk, we will overview two problems, both related to compressive sampling. The first has to do with jointly estimating all of the channel responses between an array of sources and an array of receivers. We will show that all of the sources emit random probe signals simultaneously, and the channel response between each source receiver pair is sparse, then these individual responses can be untangled by solving an L1 minimization program. The required length of the probe signals scales roughly as the joint sparsity of the channel responses.

In the second part of the talk, we will discuss recent progress on algorithms aimed at making compressive sampling ``dynamic''. We will show how the solutions to L1 optimization programs can be efficiently updated as 1) the signal we are measuring changes, and 2) new measurements are added, and stale ones are removed. The algorithms are based on homotopy methods, and are somewhat analogous to recursive least-squares in that they can be reduced to a series of low-rank updates.

.