
Compressed Sensing: Theory, Algorithms and Application
Acquisition comprimée : Théorie, algorithmes et application
(Org: **Michael Friedlander, Felix Herrmann** and/et **Ozgur Yilmaz** (UBC))

LORNE APPLEBAUM, Princeton University

Multiuser Detection in Asynchronous On–Off Random Access Channels Using Lasso

We consider on–off random access channels where users transmit either a one or a zero to a base station. Such channels represent an abstraction of control channels used for scheduling requests in third-generation cellular systems and uplinks in wireless sensor networks deployed for target detection. A key characteristic of these systems is their asynchronous nature. We will introduce a Lasso-based scheme for multiuser detection in asynchronous on–off random access channels that does not require knowledge of the delays or the instantaneous received signal-to-noise ratios of the individual users at the base station. For any fixed maximum delay in the system, the proposed scheme allows an exponential number of total users with respect to code length—achieving almost the same problem dimension scaling relationships as that required in the ideal case of fully synchronous channels. Further, the computational complexity of the proposed scheme differs from that of a similar oracle-based scheme with perfect knowledge of the user delays by at most a log factor. The results presented here are non-asymptotic, in contrast to previous work for synchronous channels that only guarantees that the probability of error at the base station goes to zero asymptotically with the number of users. Finally, we give a deterministic code construction suitable for the delay agnostic system. The code construction is based on a cyclic code in which equivalence classes are assigned to users. The code’s low coherence permits recovery guarantees with Lasso.

STEPHEN BECKER, California Institute of Technology

First-order methods for constrained linear inverse problems

Many algorithms have recently been proposed to solve the unconstrained forms of linear inverse problems, but few algorithms solve the constrained form. We show a general framework for solving constrained problems that applies to all problems of interest to compressed sensing. The technique is based on smoothing and solving the dual formulation. Using this method, it is possible to solve problems such as the Dantzig Selector, or composite problems such as minimizing a combination of the TV norm and weighted ℓ_1 norm. Additionally, we discuss recent results about exact regularization and about accelerated continuation. <http://arxiv.org/abs/1009.2065>

JIM BURKE, University of Washington

Sparsity and Nonconvex Nonsmooth Optimization

Sparsity (or parsimonious) optimization is a framework for examining the trade-off between optimality and the number independent variables to optimize over. Much of this framework was first developed for application to a range of problems in statistics where the goal is to explain as much of a data as possible with the fewest number of explanatory variables. Within the past few years, the general methodology of sparsity optimization has found its way into a wide range of applications. In this talk we consider a general sparsity optimization framework for arbitrary extended real-valued objective functions that are continuous on their essential domains. General approximation properties for the associated sparsity optimization problem are given, and a fixed point iteration for the subdifferential inclusion identifying approximate stationarity is developed. Convergence results and numerical experiments are presented.

MIKE DAVIES, University of Edinburg

Rank Aware Algorithms for Joint Sparse Recovery

This talk will revisit the sparse multiple measurement vector (MMV) problem, where the aim is to recover a set of jointly sparse multichannel vectors from incomplete measurements. This problem has received increasing interest as an extension of single

channel sparse recovery, which lies at the heart of the emerging field of compressed sensing. However, MMV approximation also has links with work on direction of arrival estimation in the field of array signal. Inspired by these links, we consider a new family of MMV algorithms based on the well-know MUSIC method in array processing highlighting the role of the rank of the unknown signal matrix in determining the difficulty of the recovery problem. We will show that the recovery conditions for our new algorithms take into account the observed rank of the signal matrix. This is in contrast to previously proposed techniques such as Simultaneous Orthogonal Matching Pursuit (SOMP) and mixed norm minimization techniques which we show to be effectively blind to the rank information. Numerical simulations demonstrate that our new rank aware techniques are significantly better than existing methods in dealing with multiple measurements.

LAURENT DEMANET, MIT

Fitting matrices from applications to random vectors

What can be determined about the inverse A^{-1} of a matrix A from one application of A to a vector of random entries? If the n -by- n inverse A^{-1} belongs to a specified linear subspace of dimension p , then come to the talk to hear which assumptions on this subspace, p , and n , guarantee an accurate recovery of A^{-1} with high probability. This randomized fitting method provides a compelling preconditioner for the wave-equation Hessian (normal operator) in seismic imaging. Joint work with Pierre-David Letourneau (Stanford) and Jiawei Chiu (MIT).

FELIX HERRMANN, the University of British Columbia

Compressive Sensing and Sparse Recovery in Exploration Seismology

During this presentation, I will talk about how recent results from compressive sensing and curvelet-based sparse recovery can be used to solve problems in exploration seismology where incomplete sampling is ubiquitous. I will also talk about how these ideas apply to dimensionality reduction of full-waveform inversion by randomly phase-encoded sources. In this second application, compressive sensing allows us to reduce the number of PDE's needed to solve this parameter-estimation problem.

GITTA KUTYNIOK, University of Osnabrueck

Geometric Separation by Single-Pass Alternating Thresholding

Modern data is customarily of multimodal nature, and analysis tasks typically require separation into the single components such as, for instance, in neurobiological imaging a separation into spines (pointlike structures) and dendrites (curvilinear structures). Although a highly ill-posed problem, inspiring empirical results show that the morphological difference of these components sometimes allows a very precise separation.

In this talk we will present a theoretical study of the separation of a distributional model situation of point- and curvilinear singularities exploiting a surprisingly simple single-pass alternating thresholding method applied to wavelets and shearlets as two complementary frames. Utilizing the fact that the coefficients are clustered geometrically in the chosen frames, we prove that at sufficiently fine scales arbitrarily precise separation is possible. Surprisingly, it turns out that the thresholding index sets even converge to the wavefront sets of the point- and curvilinear singularities in phase space and that those wavefront sets are perfectly separated by the thresholding procedure. Main ingredients of our analysis are the novel notion of cluster coherence and a microlocal analysis viewpoint.

This is joint work with David Donoho (Stanford U.).

ZHAOSONG LU, Simon Fraser University

Penalty Decomposition Methods for Rank and l_0 -Norm Minimization

In the first part of this talk, we consider general rank minimization problems. We first show that a class of matrix optimization problems can be solved as lower dimensional vector optimization problems. As a consequence, we establish that a class of rank minimization problems have closed form solutions. Using this result, we then propose penalty decomposition (PD) methods for general rank minimization problems in which each subproblem is solved by a block coordinate descend method. Under some

suitable assumptions, we show that any accumulation point of the sequence generated by our method when applied to the rank constrained minimization problem is a stationary point of a nonlinear reformulation of the problem. In the second part, we consider general l_0 -norm minimization problems. We first reformulate the l_0 -norm constrained problem as an equivalent rank minimization problem and then apply the above PD method to solve the latter problem. Further, by utilizing the special structures, we obtain a PD method that only involves vector operations. Under some suitable assumptions, we establish that any accumulation point of the sequence generated by the PD method satisfies a first-order optimality condition that is generally stronger than one natural optimality condition. Finally, we test the performance of our PD methods on matrix completion, nearest low-rank correlation matrix, sparse logistic regression and sparse inverse covariance selection problems. The computational results demonstrate that our methods generally outperform the existing methods in terms of solution quality and/or speed.

HASSAN MANSOUR, University of British Columbia

Recovery of Compressively Sampled Signals Using Partial Support Information

In this talk, we discuss the recovery conditions of weighted ℓ_1 minimization for signal reconstruction from compressed sensing measurements when partial support information is available. We show that if at least 50% of the (partial) support information is accurate, then weighted ℓ_1 minimization is stable and robust under weaker conditions than the analogous conditions for standard ℓ_1 minimization. Moreover, weighted ℓ_1 minimization provides better bounds on the reconstruction error in terms of the measurement noise and the compressibility of the signal to be recovered. We illustrate our results with extensive numerical experiments on synthetic as well as audio and video signals.

ALI PEZESHKI, Colorado State University

Sense and Sensitivity: Model Mismatch in Compressed Sensing

We study the sensitivity of compressed sensing to mismatch between the assumed and the actual models for sparsity. We start by analyzing the effect of model mismatch on the best k -term approximation error, which is central to providing exact sparse recovery guarantees. We establish achievable bounds for the ℓ_1 error of the best k -term approximation and show that these bounds grow linearly with the signal dimension and the mismatch level between the assumed and actual models for sparsity. We then derive bounds, with similar growth behavior, for the basis pursuit ℓ_1 recovery error, indicating that the sparse recovery may suffer large errors in the presence of model mismatch. Although, we present our results in the context of basis pursuit, our analysis applies to any sparse recovery principle that relies on the accuracy of best k -term approximations for its performance guarantees. We particularly highlight the problematic nature of model mismatch in Fourier imaging, where spillage from off-grid DFT components turns a sparse representation into an incompressible one. We substantiate our mathematical analysis by numerical examples that demonstrate a considerable performance degradation for image inversion from compressed sensing measurements in the presence of model mismatch.

This is joint work with Yuejie Chi, Louis Scharf, and Robert Calderbank.

HOLGER RAUHUT, Hausdorff Center for Mathematics, University of Bonn

Recovery of functions in high dimensions via compressive sensing

Compressive sensing predicts that sparse vectors can be recovered efficiently from highly undersampled measurements. It is known in particular that multivariate sparse trigonometric polynomials can be recovered from a small number of random samples. Classical methods for recovering functions in high spatial dimensions usually suffer the curse of dimension, that is, the number of samples scales exponentially in the dimension (the number of variables of the function). We introduce a new model of functions in high dimensions that uses "sparsity with respect to dimensions". More precisely, we assume that the function is very smooth in most of the variables, and is allowed to be rather rough in only a small but unknown set of variables. This translates into a certain sparsity model on the Fourier coefficients. Using techniques from compressive sensing, we are able to recover functions in this model class efficiently from a small number of samples. In particular, this number scales only logarithmically in the spatial dimension - in contrast to the exponential scaling in classical methods.

BENJAMIN RECHT, University of Wisconsin-Madison
The Convex Geometry of Inverse Problems

Building on the success of generalizing compressed sensing to matrix completion, this talk discusses progress on further extending the catalog of objects and structures that can be recovered from partial information. I will focus on a suite of data analysis algorithms designed to decompose signals into sums of atomic signals from a simple but not necessarily discrete set. These algorithms are derived in a convex optimization framework that encompasses previous methods based on ℓ_1 -norm minimization and nuclear norm minimization for recovering sparse vectors and low-rank matrices. I will discuss general recovery guarantees and implementation schemes for this suite of algorithms and will describe several example classes of atoms and applications.

JUSTIN ROMBERG, Georgia Institute of Technology
Random coding for forward modeling

Compressed sensing has shown us that sparse signal acquisition can be made efficient by injecting randomness into the measurement process. In this talk, we will show how these same ideas can be used to dramatically reduce the computation required for two types of simulation problems in acoustics. In the first, we show how all of the channels in a multiple-input multiple-output system can be acquired jointly by simultaneously exciting all of the inputs with different random waveforms, and give an immediate application to seismic forward modeling. In the second, we consider the problem of acoustic source localization in a complicated channel. We show that the amount of computation to perform “matched field processing” (matched filtering) can be reduced by precomputing the response of the channel to a small number of dense configurations of random sources.

RAYAN SAAB, The University of British Columbia
Sobolev Duals of Random Frames and Sigma-Delta Quantization for Compressed Sensing

Compressed sensing, as a signal acquisition technique, has been shown to be highly effective for dimensionality reduction. On the other hand, quantization of compressed sensing measurements has been a relatively under-addressed topic. In particular, the results of Candes, Romberg and Tao, and of Donoho guarantee that if a uniform quantizer of step size δ is used to quantize m measurements $y = \Phi x$ of a k -sparse signal $x \in \mathbb{R}^N$, where Φ satisfies the restricted isometry property, then the reconstruction error via ℓ_1 -minimization is $O(\delta)$. This is the simplest and most commonly assumed approach for quantization of compressed sensing measurements.

On the other hand, in this talk we show that if instead of uniform quantization, an r th order $\Sigma\Delta$ quantization scheme with the same output alphabet is used to quantize y , then there is an alternative recovery method via Sobolev dual frames which guarantees a reduction of the approximation error by a factor of $(m/k)^{(r-1/2)\alpha}$ for any $0 < \alpha < 1$, if $m \gtrsim_r k(\log N)^{1/(1-\alpha)}$. The result holds with high probability on the initial draw of the measurement matrix Φ from the Gaussian distribution, and uniformly for all k -sparse signals x that satisfy a mild size condition on their supports.

THOMAS STROHMER, University of California, Davis
How sparsity can enrich wireless communications

We demonstrate how concepts from compressive sensing and random matrix theory can combat some challenging problems of next-generation wireless communication systems.

EWOUT VAN DEN BERG, Stanford University
Spot - A linear operator toolbox for Matlab

Spot is a Matlab toolbox for the construction, application, and manipulation of linear operators. One of the main achievements of the package is that it provides operators in such a way that they are as easy to work with as explicit matrices, while represented

implicitly. This combines the benefits of explicit matrices with the scalability of implicit representations, thus clearing the way for fast prototyping with complex and potentially large linear operators. (This is joint work with Michael Friedlander.)

RACHEL WARD, Courant Institute, NYU

New and improved Johnson-Lindenstrauss embeddings via the Restricted Isometry Property

The Johnson-Lindenstrauss (JL) Lemma states that any set of p points in high dimensional Euclidean space can be embedded into $O(\delta^{-2} \log(p))$ dimensions, without distorting the distance between any two points by more than a factor between $1 - \delta$ and $1 + \delta$. We establish a "near-equivalence" between the JL Lemma and the Restricted Isometry Property (RIP), a well-known concept in the theory of sparse recovery often used for showing the success of ℓ_1 -minimization. In particular, we show that matrices satisfying the Restricted Isometry of optimal order, with randomized column signs, provide optimal Johnson-Lindenstrauss embeddings up to a logarithmic factor in N . Our results have implications for dimensionality reduction and sparse recovery: on the one hand, we arrive at the best known bounds on the necessary JL embedding dimension for a wide class of structured random matrices; on the other hand, our results expose several new families of universal encoding strategies in compressed sensing.

This is joint work with Felix Krahmer.

REBECCA WILLET, Duke University

Compressed Sensing with Poisson Noise

Compressed sensing has profound implications for the design of new imaging and network systems, particularly when physical and economic limitations require that these systems be as small and inexpensive as possible. However, several aspects of compressed sensing theory are inapplicable to real-world systems in which noise is signal-dependent and unbounded. In this work we discuss some of the key theoretical challenges associated with the application of compressed sensing to practical hardware systems and develop performance bounds for compressed sensing in the presence of Poisson noise. We develop two novel sensing paradigms, based on either pseudo-random dense sensing matrices or expander graphs, which satisfy physical feasibility constraints. In these settings, as the overall intensity of the underlying signal increases, an upper bound on the reconstruction error decays at an appropriate rate (depending on the compressibility of the signal), but for a fixed signal intensity, the error bound actually grows with the number of measurements or sensors. This surprising fact is both proved theoretically and justified based on physical intuition.