News & Events

IDeAS Seminar

Thursday;   3:45 pm, 102A McDonnell Hall Seminar Archives
collapse all – expand all
To speak at this seminar, contact Afonso Bandeira.

Spring 2013 Collapse/Expand

Me

Date: January 31

Room: Fine 214

Speaker: Jiawei Chiu

Title: Towards practical sparse Fourier transform

Abstract: Suppose we have a signal of size N that is S-sparse in frequency space. We seek algorithms that can compute its Fourier transform faster than the FFT. There are some existing algorithms that run in O(S) time ignoring log factors, where S is the sparsity of the signal in frequency space. However, for a fixed N=2^22, these algorithms beat FFTW only when S is roughly less than 200. In this work, we extend sFFT3, a non-robust algorithm, to a robust O(S)-time algorithm that runs at least 5 times faster than existing O(S)-time algorithms. The main new ingredient is the use of the matrix pencil method to detect so-called mode collisions. If time permits, we will discuss how matrix skeletons or CUR representations can be computed in time independent of the size of the matrix.

Date: February 5

Room: Fine 214

Speaker: Irene Waldspurger, Yale University

Title: Phase retrieval for the wavelet transform

Abstract: The modulus of the wavelet transform is a common way to represent audio signals. We are interested in the corresponding inverse problem : can a signal be reconstructed from the modulus of its wavelet transform ? For carefully chosen wavelets, this problem amounts to finding the zeros of an holomorphic function. In this case, the wavelet transform's modulus uniquely determines the function, up to a global phase. Moreover, reconstruction is continuous. However, as continuity is not uniform, there are instabilities. We will describe two types of instabilities, which correspond to audio signals that our ear cannot distinguish. The holomorphic structure of the problem leads to think that all instabilities are combinations of these two types. Finally, we will present a reconstruction algorithm based on convexification techniques. This algorithm is general and may be applied to any phase retrieval problem. Numerical tests show that it provides accurate and stable results.

Date: February 21

Room: 102A McDonnell Hall

Speaker: Kunal Chaudhury, Princeton University

Title: Global registration of multiple point clouds using semidefinite programming

Abstract: We consider the problem of positioning a massive point cloud P from the positions (possibly noisy) of smaller point clouds P_1,...,P_M. This problem comes up in distributed approaches to sensor network localization and molecular conformation, and also in certain computer vision applications. The model we consider is that the points in each P_i are obtained from the corresponding points in P through a single rigid transform (model variable). In this talk, we will formulate a global optimization framework for this model, and show how the variables (rigid transforms) can be efficiently optimized using semidefinite relaxation. This is a generalization of the well-known SVD procedure [Horn, Arun, et al.] for registering two point clouds. We will compare the proposed method with existing methods in terms of complexity and performance. If time permits, we will also present some results on exact recovery and stability for the network localization problem. This is a joint work with Amit Singer and Yuehaw Khoo.

Date: February 28

Room: 102A McDonnell Hall

Speaker: Tong Zhang, Rutgers University

Title: Stochastic Dual Coordinate Ascent and its Proximal Extension for Regularized Loss Minimization

Abstract: Stochastic Gradient Descent (SGD) has become popular for solving large scale supervised machine learning optimization problems such as SVM, due to their strong theoretical guarantees. While the closely related Dual Coordinate Ascent (DCA) method has been implemented in various software packages, it has so far lacked good convergence analysis. We present a new analysis of Stochastic Dual Coordinate Ascent (SDCA) showing that this class of methods enjoy strong theoretical guarantees that are comparable or better than SGD. This analysis justifies the effectiveness of SDCA for practical applications. Moreover, we introduce a proximal version of dual coordinate ascent method. We demonstrate how the derived algorithmic framework can be used for numerous regularized loss minimization problems, including L1 regularization and structured output SVM. The convergence rates we obtain match or improve state-of-the-art results. Joint work with Shai Shalev-Shwartz

---

Date: March 7

Room: 102A McDonnell Hall

Speaker: Bart Vandereycken

Title: Rank-structured tensors in numerical analysis using optimization on manifolds.

Abstract: Approximating high-dimensional functions is in general intractable, even for the class of infinitely differentiable multivariate functions. However, certain functions, like solutions of elliptic high-dimensional PDEs, have sufficiently high regularity that they have a tractable and data-sparse representation. Sparse grids are a well-established technique for doing this, but in this talk I will focus on an alternative approach, namely, that of rank-structured tensors. Not to be confused with the more generally known tensor rank--which is NP hard--the rank-structured tensors in this talk lead to a stable and computationally efficient format. They were only recently introduced in the numerical analysis community by [Hackbusch-Kuehn, 2009] and [Oseledets-Tyrtyshnikov, 2009], but similar formats have already a longstanding tradition in computational chemistry to simulate quantum systems [White, 1992]. I will give an overview of the so-called hierarchical Tucker (HT) format and its application for a typical problem in numerical analysis: the solution of a high-dimensional, possibly time-dependent PDE directly in a data-sparse represention. The approach is conceptually simple: constrain the original problem to the manifold of fixed-rank HT tensors and subsequently solve and analyze this problem using techniques from optimization on manifolds. Where possible, I will show numerical results to compare with alternative approaches, like alternating least squares. Parts of this are based on joint work with D. Kressner, Ch. Lubich, T. Rohwedder, R. Schneider, and M. Steinlechner.

Date: March 14

Room: 102A McDonnell Hall

Speaker: Ankur Moitra, Institute for Advanced Study

Title: Provable Algorithms for Nonnegative Matrix Factorization and Beyond

Abstract: The nonnegative matrix factorization problem has important applications throughout machine learning where it is used to uncover latent statistical relationships present in data, that can then be used for clustering, information retrieval, recommendation systems etc. As is often the case, this problem is NP-hard when considered in full generality. However, we introduce a sub-case called separable nonnegative matrix factorization that we believe is the right notion in various contexts. We give a polynomial time algorithm for this problem, and leverage this algorithm to efficiently learn the topics in a Latent Dirichlet Allocation model and various other topic models. In fact, these algorithms are not only interesting from a theoretical standpoint but in fact run orders of magnitude faster than the existing best algorithms for these tasks, without sacrificing and in many cases improving the quality of the output. There are many natural questions about how these approaches can be extended to more general settings, and whether these algorithms can be successfully applied to even larger data sets than we have had the tools to explore thus far. Joint work with Sanjeev Arora, Rong Ge and Ravi Kannan.

Date: March 28

Room: 102A McDonnell Hall, **3:00 p.m. (special start time!)**

Speaker: Teng Zhang

Title: Robust subspace recovery by Tyler's M-estimator

Abstract: We show that Tyler's M-estimator has a nice property for robust subspace recovery (though it was mainly used for robust covariance estimation). Specifically, when inliers are sampled from a subspace and the percentage of inliers is larger than some threshold, then the underlying subspace can be exactly recovered by Tyler's M-estimator. The main tools for the proof are the geodesic convexity of the objective function and the majorization-minimization property of the associated algorithm. Besides, we will also discuss the stability of Tyler's M-estimator, and show similar convexity to other maximum-likelihood estimation of covariance matrix.

Date: April 4

Room: Jadwin A09

Speaker: Andy Zhu

Title: Multireference Alignment via Convex Optimization

Abstract: The multireference alignment problem consists of estimating a signal from noisy shifted observations of it. If the shifts were known one could simply shift back each observation, and then average to diminish the noise. However, in relevant applications, the shifts are unknown. We present several methods to estimate the unknown shifts and the signal. Our main contribution is an poly-time approximation algorithm to solve this problem inspired by a certain semidefinite programming based approach to the Unique Games problem, which seems to have comparable performance to the (computationally hard) maximum likelihood estimator. We also mention how we can leverage recent results about symmetry reduction in semidefinite programs from representation theory. Joint work with A. S. Bandeira, M. Charikar, and A. Singer.

Date: April 10

Room: 102A McDonnell Hall

Speaker: Matan Gavish, Stanford University

Title: Minimax Risk of Matrix Denoising by Singular Value Thresholding

Abstract: In matrix denoising, we estimate an unknown real m-by-n matrix X from a single noisy measurement Y=X+Z, where Z is assumed to have i.i.d normal entries. A popular matrix denoising scheme is Singular Value Thresholding (SVT), which applies soft thresholding (with a common threshold) to each of the singular values of the data matrix Y. I'll show an explicit formula for the minimax MSE of SVT over matrices of rank <= r, and the corresponding optimal location for the threshold of SVT. I will also compare the minimax MSE of SVT to the minimax MSE over all possible matrix denoisers. For example, for n=m, we will see that SVT is no more than a factor of 3 suboptimal. Matrix denoising is connected to compressed sensing of matrices: there is empirical evidence that the minimax formula I'll present accurately predicts the phase transition of matrix recovery from Gaussian measurements. Joint work with David Donoho.

Date: April 18

Room: 102A McDonnell Hall

Speaker: Nicolas Boumal

Title: Relaxing is not the end of it

Abstract: Synchronization of rotations is the problem of estimating a set of rotations from measurements of pairwise relative rotations. Several convex relaxations of this problem have been proposed. These relaxations enjoy excellent theoretical performance guarantees. The relaxation approaches involve a final projection step, that brings back the solution of the relaxed problem to the original feasible set. This projection step though does not, in general, provide even a local

optimizer of the original problem. In this presentation, I will argue that relaxing is not the end of it: further reaching for a critical point of the problem we actually want to solve can make a big difference. Indeed, the algorithm we propose appears to be efficient, as compared to fundamental Cramér-Rao bounds. If time allows, I will present Manopt, a matlab toolbox for optimization on manifolds. This is the toolbox we use to perform the nonlinear optimization step that follows the projection step in synchronization. Such a toolbox is of interest to refine relaxed solutions in a wide range of applications. Joint work with Amit Singer and Pierre-Antoine Absil.

Date: April 25

Room: 102A McDonnell Hall

Speaker: Emmanuel Abbe, Princeton University

Title: Recent developments in low-complexity coding

Abstract: I will start with a basic question in coding theory: how to compress biased bits with a linear map. After revisiting the fundamental result of Shannon, I will discuss the challenge of achieving the limit with low-complexity matrices, and present two results using sparse and polar code matrices. I will then extend the results to multiple sources, and propose a framework for compression over the reals. This will take us to the problem of sensing high-dimensional real signals.

Date: May 2

Room: 102A McDonnell Hall

Speaker: Onur Ozyesil

Title: TBA

Abstract: TBA

Fall 2012 Collapse/Expand

Date: September 19

Room: Fine 214

Speaker: Lanhui Wang, PACM

Title: Robust Synchronization and Exact Recovery of Rotations

Abstract: The synchronization problem over the special orthogonal group $SO(d)$ consists of estimating a set of unknown rotations $R_1,R_2,\ldots,R_n$ from noisy measurements of a subset of their pairwise ratios $R_{i}^{-1}R_{j}$. The problem has found applications in computer vision, computer graphics, and sensor network localization, among others. Its least squares solution can be approximated by either spectral relaxation or semidefinite programming followed by a rounding procedure, analogous to the approximation algorithms of \textsc{Max-Cut}. The contribution of this talk is three-fold: First, we introduce a robust penalty function involving the sum of unsquared deviations and derive a relaxation that leads to a convex optimization problem; Second, we show that the alternating direction method converges linearly to the unique minimizer; Finally, we prove that the rotations are exactly recovered for a certain model for the measurement noise and for the measurement graph, exhibiting a phase transition behavior in terms of the proportion of noisy measurements. Numerical simulations confirm the phase transition behavior for our method as well as its improved accuracy compared to existing methods.

Date: October 3

Room: Fine 214

Speaker: Nicolas Boumal, U.C. Louvain in Belgium

Title: Cramér-Rao bounds for synchronization of rotations (are structured by the pseudoinverse of the Laplacian)

Abstract: Synchronization of rotations is the problem of estimating a set of rotations $R_1, ..., R_N$ in $SO(n)$ based on noisy measurements of relative rotations $R_i R_j^T$. This fundamental problem has found many recent applications, most importantly in structural biology. We provide a framework to study synchronization as estimation on Riemannian manifolds. As a main contribution, we provide the Cramér-Rao bounds for synchronization. We find that these bounds are structured by the pseudoinverse of the measurement graph Laplacian, where edge weights are proportional to measurement quality. We leverage this to provide interpretation and visualization tools for these bounds.

Date: November 27

Room: Fine 214 **(special time - 4:30 p.m.)**

Speaker: Ohad Medalia, Zurich University

Title: Visualizing cellular processes at the molecular level by cryo-electron tomography

Abstract: Visualization of the three-dimensional (3-D) organization of a eukaryotic cell, with its dynamic organelles, cytoskeletal structures, and distinct protein complexes in their native context, requires a non-invasive imaging technique of high resolution combined with a method of arresting cellular elements in their momentary state of function. Vitrification of cells ensures close-to-life preservation of the molecular architecture of actin networks and organelles. With the advent of automated electron tomography it has becomes possible to obtain tomographic data sets of frozen hydrated specimen. By electron tomography 3-D information from large pleomorphous structures, as cell organelles or whole cells can be retrieved with 'molecular resolution'. At that resolution it becomes possible to detect and identify specific macromolecular complexes on the basis of their structural signature. Here we employed cryo-electron tomography in order to reconstruct macromolecular assemblies from Eukaryotic cells. The possibilities, limitations and needs for novel algorithms for analysing the data, will be discussed.

Date: November 28

Room: Fine 214

Speaker: Joan Bruna, Courant Institute

Title: Scattering Representations for Recognition

Abstract: Pattern and Texture Recognition are high level tasks requiring signal representations with specific invariance, stability and consistency properties, which are not satisfied by linear decompositions. Scattering operators cascade wavelet decompositions and complex modulus, followed by a lowpass filtering. They define a non-linear representation which is locally translation invariant and Lipschitz continuous to the action of diffeomorphisms. They also define a texture representation capturing high order moments and which can be consistently estimated from few realizations. This talk will present the main mathematical properties of scattering operators and cover their main applications on pattern and texture recognition. We will put special emphasis on the scattering representation of multifractal measures, which bring a novel fractal descriptor and yield stable, consistent estimates for fractal characterization.

Date: December 12

Room: Fine 214

Speaker: Dustin G. Mixon, Air Force Institute of Technology (AFIT)

Title: Phase retrieval: Approaching the theoretical limits in practice

Abstract: In many areas of imaging science, it is difficult to measure the phase of linear measurements. As such, one often wishes to reconstruct a signal from intensity measurements, that is, perform phase retrieval. Today, very little is known about how to design injective intensity measurements, let alone stable measurements with efficient reconstruction algorithms. In fact, the state of the art applies certain routines called PhaseLift or PhaseCut, but for these, performance guarantees are only available when the measurement vectors are Gaussian random, which cannot be used in many applications. This talk will help fill the void - I will discuss a wide variety of recent results in phase retrieval, including various conditions for injectivity and stability (joint work with Afonso S. Bandeira (Princeton), Jameson Cahill (U Missouri) and Aaron A. Nelson (AFIT)) as well as measurement designs based on spectral graph theory which allow for efficient reconstruction (joint work with Boris Alexeev (Princeton), Afonso S. Bandeira (Princeton) and Matthew Fickus (AFIT)). In particular, I will show how Fourier-type tricks can be leveraged in concert with this graph-theoretic design to produce quasi-random aperatures for coherent diffractive imaging (joint work with Afonso S. Bandeira (Princeton) and Yutong Chen (Princeton).

Spring 2012 Collapse/Expand

return to top

Date: February 8

Room: Fine 214

Speaker: Lanhui Wang, PACM

Title: A Fourier-based Approach for Iterative 3D Reconstruction from Cryo-EM Images

Abstract: A major challenge in single particle reconstruction methods using cryo-electron microscopy is to attain a resolution sufficient to interpret fine details in three-dimensional (3D) macromolecular structures. Obtaining high resolution 3D reconstructions is difficult due to unknown orientations and positions of the imaged particles, possible incomplete coverage of the viewing directions, high level of noise in the projection images, and limiting effects of the contrast transfer function of the electron microscope. In this paper, we focus on the 3D reconstruction problem from projection images assuming an existing estimate for their orientations and positions. We propose a fast and accurate Fourier-based Iterative Reconstruction Method (FIRM) that exploits the Toeplitz structure of the operator ${\bf A}^{*}{\bf A}$, where $\bf A$ is the forward projector and ${\bf A}^{*}$ is the back projector. The operator ${\bf A}^{*}{\bf A}$ is equivalent to a convolution with a kernel. The kernel is pre-computed using the non-uniform Fast Fourier Transform and is efficiently applied in each iteration step. The iterations by FIRM are therefore considerably faster than those of traditional iterative algebraic approaches, while maintaining the same accuracy even when the viewing directions are unevenly distributed. The time complexity of FIRM is comparable to the gridding-based direct Fourier inversion method. Moreover, FIRM combines images from different defocus groups simultaneously and can handle a wide range of regularization terms. We provide experimental results on simulated data that demonstrate the speed and accuracy of FIRM in comparison with current methods.

Date: February 15

Room: Fine 214

Speaker: Jane Zhao, PACM

Title: Rotationally Invariant Image Representation for Viewing Angle Classification in Cryo-EM

Abstract: We introduce a new rotationally invariant viewing angle classification method for identifying, among a large number of Cryo-EM projection images, similar views without prior knowledge of the molecule. Our rotationally invariant features are based on the bispectrum. Each image is first expanded in an orthonormal steerable basis to generate expansion coefficients. Rotating an image is equivalent to phase shifting the expansion coefficients. Thus we are able to extend the theory of bispectrum of 1D periodic signals to 2D images. The randomized PCA algorithm is then used to efficiently reduce the dimensionality of the bispectrum coefficients, enabling fast computation of the similarity between any pair of images. The nearest neighbors provide an initial classification of similar viewing angles. In this way, rotational alignment is only performed for images with their nearest neighbors. The initial nearest neighbor classification and alignment are further improved by a new classification method called vector diffusion map.

Date: February 22

Room: Fine 214

Speaker: Nikhil Srivastava, Computer Science, Princeton

Title: Covariance Estimation for Distributions with 2+εMoments

Abstract: We study the minimal sample size N=N(n) that suffices to estimate the covariance matrix of an n-dimensional distribution by the sample covariance matrix in the operator norm, and with an arbitrary fixed accuracy. We establish the optimal bound N = O(n) for every distribution whose k-dimensional marginals have uniformly bounded 2+\epsilon moments outside the sphere of radius O(\sqrt{k}). In the specific case of log-concave distributions, this result provides an alternative approach to the Kannan-Lovasz-Simonovits problem, which was recently solved by Adamczak, Litvak, Pajor and Tomczak-Jaegermann. Moreover, a lower estimate on the covariance matrix holds under a weaker assumption -- uniformly bounded 2+\epsilon moments of one-dimensional marginals. Our argument proceeds by randomizing the spectral sparsification technique of Batson, Spielman and Srivastava. The spectral edges of the sample covariance matrix are controlled via the Stieltjes transform evaluated at carefully chosen random points.

Date: March 7

Room: Fine 214

Speaker: George F. Young, Mechanical & Aerospace Engineering, Princeton University

Title: Optimising Robustness of Consensus to Noise

Abstract: One of the key benefits of consensus in a multi-agent system is that it can improve decision making in the presence of uncertainty. However, the act of communication itself will usually introduce additional noise to the system. For this reason it is important to understand how well a group can reach and maintain consensus in the presence of communication noise and in particular whether it is possible to minimise the disturbances created by this noise. Here we use a measure of group dispersion due to noise, which can be interpreted as the H2 norm of the consensus system, to relate the communication graph between the agents to the group robustness. We explore some fundamental bounds on performance and investigate how to improve robustness when the graph is a tree. Our metric is also used to evaluate different sensing strategies in starling flocks to determine which ones achieve robust consensus in the most efficient manner.

Date: March 14

Room: Fine 214

Speaker: Hau-Tieng Wu, PACM, Princeton University

Title: Manifold adaptive local linear regression and its application to sleep depth estimation

Abstract: We study nonparametric regression with high-dimensional massive dataset, when the predictors lie on an unknown, lower-dimensional geometric object. In particular, we focus on the sleep depth estimation problem, where the dataset is generated by applying Synchrosqueezing transform, a time frequency analysis technique, to the sleep data. When the geometric object is a smooth manifold, our approach to the nonparametric regression is to reduce the dimensionality first and then construct the local linear regression (LLR) directly on the tangent plane approximation to the manifold. Under mild conditions, asymptotic expressions for the conditional mean squared error of the proposed estimator are derived for both the interior and the boundary cases. One implication of these results is that the optimal convergence rate depends only on the intrinsic dimension d of the manifold, but not on the ambient space dimension p. Another implication is that the estimator is design adaptive and automatically adapts to the boundary of the unknown manifold. The proposed method has a strong connection with manifold learning and the second implication leads to a new diffusion map framework.

Date: March 28

Room: Fine 214

Speaker: Dimitris Giannankis, Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences

Title: Capturing intermittent and low-frequency variability in high-dimensional time series through nonlinear Laplacian spectral analysis

Abstract: Nonlinear Laplacian spectral analysis (NLSA) is a method for spatiotemporal analysis of high-dimensional data, which represents temporal patterns via orthonormal basis functions on the nonlinear data manifold. Through the use of such basis functions (determined by means of graph Laplace-Beltrami eigenfunction algorithms), NLSA captures intermittency, rare events, and other nonlinear dynamical features which are not accessible through classical linear approaches such as singular spectrum analysis. We discuss applications of NLSA to detection of decadal and intermittent variability in the North Pacific sector of comprehensive climate models, and dimension reduction of a chaotic low-order model of the atmosphere.

Date: April 11

Room: Fine 214

Speaker: Xiuyuan Cheng - PACM Graduate Student

Title: The Spectrum of Random Kernel Matrices

Abstract: This talk is about random matrix problems motivated by the use of kernel method in high-dimensional statistical learning. We will present a mathematical result concerning the limiting spectral density of inner-product kernel matrices built from p-dimensional Gaussian vectors X_1, ... X_n, where both p and n increase to infinity with p/n staying at a constant. The (i,j)-th entry of the random matrix equals f(X_i^TX_j), where f is a real-valued function. As a well-known classical result, the limiting spectral density is the Marcenko-Pastur distribution when f is a linear function. We will show that, for a large class of kernel function f, the limiting spectral density exists and is dictated by a cubic equation involving its Stieltjes transform. The new family of limiting densities includes the Marcenko-Pastur distribution and Wigner’s semi-circle as special cases. We will also talk about a "spiking" model of the kernel matrices where the vector X_i's admit a model of "low-dimensional information corrupted by high-dimensional noise".

Date: April 18

Room: Fine 214

Speaker: Afonso Bandeira - PACM Graduate Student

Title: Cheeger Inequality for the Graph Connection Laplacian

Abstract: The $O(d)$ Synchronization problem consists of estimating a set of unknown orthogonal transformations O_i from noisy measurements of a subset of the pairwise ratios $O_iO_j^{-1}$. We formulate and prove a Cheeger-type inequality that relates a measure of how well it is possible to solve the $O(d)$ synchronization problem with the spectra of an operator, the graph Connection Laplacian. We also show how this inequality provides a worst case performance guarantee for a spectral method to solve this problem (This is joint work with Amit Singer (Princeton) and Daniel Spielman (Yale)). These guarantees also play a major role on showing the stability of a certain method to solve the problem of reconstruction without phase (this part is joint work with: Boris Alexeev (Princeton), Dustin Mixon (Princeton) and Matthew Fickus (Air Force Inst. Tech.)).

Date: April 19

Room: Joseph Henry Room - Jadwin Hall

Speaker: Michael McCoy, California Institute of Technology

Title: Random geometry meets data analysis: Sharp recovery bounds for convex deconvolution

Abstract: Consider the common situation where observed data consists of the superposition of two signals. Some examples include: an image of the night sky containing both stars and galaxies; a communications message with impulsive noise; and a low rank matrix obscured by sparse corruptions. Deconvolution is the problem of determining the constituent signals from their superposition. A fundamental question is "When is deconvolution possible with a tractable algorithm?" We describe a convex optimization framework for deconvolution, and provide a geometric characterization of success in this framework. This geometric viewpoint, coupled with a natural incoherence model, leads us into the realm of random geometry. A powerful result from spherical integral geometry yields an exact formula for the probability that our program succeeds. This formula leads to simple summary statistics that reveal sharp phase transitions between success and failure, and important theoretical properties of convex regularizers. We apply our results to deconvolving the superposition of sparse vectors in random bases, a stylized robust communications protocol, and determining a low rank matrix corrupted by a matrix that is sparse in a random basis. We show that empirical results closely match our theoretical bounds. Joint work with Joel Tropp

---

Date: April 24

Room: 214 Fine Hall

Speaker: Ronen Talmon, Yale University

Title: Differential Stochastic Sensing: Intrinsic Modeling of Random Time Series with Applications to Nonlinear Tracking

Abstract: Many natural and artificial high-dimensional data sets are controlled by few lower-dimensional factors or drivers. As a result, the data is often highly structured and does not fill uniformly the high-dimensional space. In this talk, we present a "differential stochastic sensing" framework for inferring the independent controlling factors (or drivers) of high-dimensional time series. This approach provides intrinsic global modeling for noisy observations based on anisotropic diffusion and local dynamical models. The idea is to implicitly solve local differential equations based on local density estimates in a global graph-based mechanism that inverts the observation function and reveals the underlying structure.

Moreover, it implicitly recovers the dynamical model of the data. Hence, it provides a foundation for sequential processing that is applied to nonlinear tracking problems. We revisit classical Bayesian filtering methods and discuss their relationship to the proposed approach. In addition, we show that the proposed intrinsic modeling is invariant under different observation schemes and is noise resilient. Hence, it may be applied to a wide variety of applications. In this talk, we demonstrate applications to the processing of financial and neuroscience time series, and biological and medical imaging.

Date: April 25

Room: Fine 214

Speaker: Sewoong Oh, Massachusetts Institute of Technology

Title: Message-passing algorithms for approximate singular vector computation

Abstract: Low-rank matrix approximation is important in many applications for capturing the important aspects of data naturally described in a matrix form. In particular, we are interested in solving an inference problem using the leading singular vectors of a data matrix, which come from crowdsourcing platforms like Mechanical Turk. Crowdsourcing systems, in which numerous tasks are electronically distributed to numerous "information piece-workers", have emerged as an effective paradigm for human-powered solving of large scale problems. Because these low-paid workers can be unreliable, we need to devise schemes to infer the correct answers to these crowdsourcing tasks from possibly incorrect responses from the workers. In this talk, to solve this inference problem, we introduce a new message-passing algorithm and prove that this algorithm is asymptotically optimal through comparison to an oracle that knows the reliability of every worker. This algorithm is inspired by the power iteration method for computing the leading singular vectors, and there is an interesting relation between the fixed point of the message-passing algorithm and the leading singular vector. The extrinsic nature of message-passing allows us prove sharp asymptotic bounds on the performance using density evolution. However, tracking the densities of real-valued messages is an a priori difficult task. We establish that the messages are sub-Gaussian using recursion, and compute an upper bound on the parameters in a closed form.

Date: May 2

Room: Fine 214

Speaker: Onur Ozyesil, Princeton University

Title: Synchronization in Non-compact Groups: SE(k) Case

Abstract: We present new algorithms for the synchronization problem in SE(k), and focus on the "Compactification by Contraction" approach with applications to synthetic data measurements and to the "Sensor Network Localization" problem. We also give a performance analysis, based on random matrix theory, and identify possible generalizations of our approach for the synchronization problem on Cartan motion groups.

Date: May 24

Room: Joseph Henry Room - Jadwin Hall

Speaker: Ewout van den Berg, Stanford University

Title: Exploiting sparsity and low-dimensional structure: Two new applications

Abstract: The success of compressed sensing in reconstructing sparse signals from a small number of linear measurements has had a tremendous impact on the signal processing community. Signal sparsity is now routinely used and extensions to different types of low-dimensional structure, such a low-rank matrices, have been successfully applied and analyzed. In this talk I will present two applications whose inherent low-dimensionality can be leveraged to our advantage. The first application concerns the design of an efficient high-resolution silicon photomultiplier chip. By exploiting the spatio-temporal sparsity in photon arrival, the proposed design achieves a high level of sensitivity while maintaining a low circuit complexity. The second application concerns volume reconstruction from tilt-series data -- one of the central problems in cryo-electron tomography. In the case of a heterogeneous data set the challenge is to determine not only the relative orientation of the objects, but also their class. We pose this problem as a parameterized low-rank matrix recovery problem, which is then solved using techniques from convex optimization.

2011 Collapse/Expand

return to top

Date: February 8
Room: The Joseph Henry Room - Jadwin Hall
Speaker: Teng Zhang, University of Minnesota
Title: Modeling Data by Multiple Subspaces: Theory and Algorithms
Abstract: click to view
We study the problem of modeling data by several affine subspaces, which generalizes the common modeling by a single subspace. It arises, for example, in object tracking and structure from motion. One of the simplest methods for such modeling is based on energy minimization, where the energy involves p-th powers of distances of data points from subspaces. We first analyze under certain assumptions (e.g., spherically symmetric outliers) the effectiveness of such energy minimization for recovering all subspaces simultaneously and also recovering the most significant subspace. We reveal the following phase transition in our setting: when p<=1 the underlying subspaces can be recovered by such energy minimization; whereas when p>1 the underlying subspaces are sufficiently far from the global minimizer. Nevertheless, for more general settings (i.e., outliers which are not spherically symmetric) we can point at some disadvantages of the energy minimization strategy. In order to practically solve the problem, we present a simple and fast geometric method for multiple subspaces modeling. It forms a collection of local best fit affine subspaces, where the size of the local neighborhoods is determined automatically by the Peter Jones' beta numbers. This collection of subspaces can then be further processed in various ways. For example, greedy selection procedure according to an appropriate energy or a spectral method to generate the final model. We demonstrate the state of the art accuracy and speed of the suggested procedure on applications for several applications.

Date: February 10 - **11 a.m. (special time!)**
Room: The Joseph Henry Room - Jadwin Hall
Speaker: Yuan Yao, Peking University
Title: Topological and Geometric Methods for Data Analysis
Abstract: click to view
Over the last two decades, the world has witnessed an enormous growth in data sets that are complex, high-dimensional, and massive. Traditional techniques for analyzing data have become inadequate. In this talk we will discuss how some classical mathematics created for completely different purposes (e.g. algebraic and differential topology, combinatorial differential geometry), could nonetheless provide powerful new tools for exploring these modern data sets. In particular we will focus on some novel applications of Hodge Theory, a bridge over topology and geometry, in statistical data analysis which have never been exploited in such a perspective.

Date: April 12
Room: 138 Lewis Library
Speaker: Nicolas Boumal, Université catholique de Louvain (Belgium)
Title: Low-rank matrix completion: optimization on manifolds at work
Abstract: click to view
We consider moderately large matrices (millions or billions of entries) of low rank. We address the problem of recovering such matrices when most of the entries are unknown. Matrix completion finds applications in recommender systems. In this setting, the rows of the matrix may correspond to items and the columns may correspond to users. The known entries are the ratings given by users to some items. The aim is to predict the unobserved ratings. This problem is commonly stated in a constrained optimization framework. We follow an approach that exploits the geometry of the low-rank constraint to recast the problem as an unconstrained optimization problem on the Grassmann manifold. We then apply a superlinear Riemannian trust-region method to solve it. We improve on key aspects of existing methods such as Admira, FPCA, OptSpace, SET and others. I will demonstrate the performance of our algorithm in terms of accuracy and speed on synthetic data, and discuss real data in the conclusions.

Date: April 26
Room: 314 Fine Hall, 3:30 p.m **(Special Time and Place!)**

Speaker: Dr. Mou-Hsiung (Harry) Chang, Mathematical Sciences Division, U.S. Army Research Office
Title: Asymptotic Behaviors of Quantum Markov Semigroups
Abstract: click to view
In this talk, we discuss some large time behaviors of a quantum Markov semigroup of bounded linear transformations acting on a von Neumann algebra on a complex Hilbert space. Large deviation principle, Lyapunov stability as well as necessary and sufficient conditions for existence of an invariant state are established using integral representation of normal state. The results are directly applicable to open quantum systems.

Date: May 17
Room: 138 Lewis Library
Speaker: Dr. Hiba Abdullah, Fourier Institute
Title: Diffusion process on a flow of Riemannian manifolds
Abstract: click to view
In this talk, we create links between the properties of diffusion of the Riemannian manifold and its geometry. We embedd a family of Riemannian manifolds whose metric is time dependent, into a Hilbert space with its diffusion properties. Namely, via the eigenfunctions of the corresponding laplacian or its heat kernel. We prove that we can construct embeddings via a finite number of eigenfunctions for all families of Riemannian manifolds (M,g(t)) such that g(t) is analytic in t. Then, we construct the fundamental solution P for the non-linear heat equation acting on (M,g(t)), such that the volume (M,g(t)) is constant, and we prove that, under certain constraints, we can embed (M,g(t)) into a Hilbert space via P.

Date: September 20
Room: The Joseph Henry Room - Jadwin Hall
Speaker: Daniel Kaslovsky, Univ. of Colorado - Boulder
Title: Optimal Tangent Plane Recovery from Noisy Manifold Samples
Abstract: click to view
Efficient data processing algorithms may be realized by exploiting the low-dimensional manifold structure often inherent in a data set. We seek efficient parameterizations of such data sets via projection into appropriate manifold tangent planes. Parameterizing a data set thus becomes a problem of estimating local tangent planes from noisy manifold samples. Principal component analysis (PCA) is often the tool of choice, as it returns an optimal basis in the case of noise-free samples from a linear subspace. To process noisy data, PCA must be applied locally, at a scale small enough such that the manifold is approximately linear, but at a scale large enough such that structure may be discerned from noise. We present our approach that uses the geometry of the data to guide our definition of locality, discovering the optimal balance of this noise-curvature trade-off. Using perturbation theory of eigenspaces, we study the stability of the subspace estimated by PCA as a function of scale, and bound (with high probability) the angle it forms with the true tangent space. By adaptively selecting the scale that minimizes this bound, our analysis reveals the optimal scale for local tangent plane recovery. Applications are discussed, with a focus on image processing.
