FOCM 2014 - Workshop A2
Computational Harmonic Analysis, Image and Signal Processing

A2 - December 11, 14:35 – 15:25

Streaming signal reconstruction from generalized measurements

Justin Romberg
Georgia Institute of Technology, USA
jrom@ece.gatech.edu

The central theme of this talk is reconstructing a signal from a stream of generalized samples. This problem has a long history in the signal processing literature. To date, the majority of the results revolve around systems which take samples of and reconstruct the signal using filterbanks with multiple channels, or reconstruct the signal in "batch mode" by collecting a large number of measurements and then perform the inversion of the entire signal by solving a system of linear equations. In the first part of this talk, we will present a method for reconstructing a signal online that lies in between these two approaches. We set the reconstruction up as a linear inverse problem, and then show how to solve the system in an "online" manner.

In the second part of the talk, we show how these ideas can be extended to sparse reconstruction, where we are solving an l1-regularized inverse problem. We present a collection of homotopy-based algorithms that dynamically update the solution of the underlying L1 problem as the system changes.

Joint work with M. Salman Asif.

A2 - December 11, 15:30 – 15:55

On exact recovery of signals from the projection onto polynomial spaces

Shai Dekel
GE Global Research and School of Mathematical Sciences, Tel-Aviv, Israel
shai.dekel@ge.com

In the talk we will review some recent contributions (as well as ours) to the following prototype problem: We are given the projection of a superposition of Diracs onto a finite dimensional polynomial space over a manifold (e.g. trigonometric polynomials, algebraic polynomials, spherical harmonics) and we wish to recover the signal exactly and in particular, the locations of the knots. We will show that under a separation condition on the support of the unknown signal, there exists a unique solution through TV minimization over the space of Borel measures. Time allowing, we will present extensions to recovery of splines, streams of pulses, numerical algorithms, experimental results, stability under noise and more.

Joint work with Tamir Bendory (Technion) and Arie Feuer (Technion).

A2 - December 11, 16:00 – 16:25

Least squares regularized or constrained by L_0: relationship between their optimal solutions and properties

Mila Nikolova

CMLA – CNRS ENS-Cachan, France

nikolova@cmla.ens-cachan.fr

When looking for a sparse solution of an under-determined linear system, two popular options are to find the optimal solution of the least squares regularized by L_0 pseudo-norm using a trade-off parameter β or constrained by L_0 (known also as the K-sparsity constrained problem).

We provide important features of the optimal solutions of both problems. We analyse in depth the relationship between the sets of optimal solutions of these two nonconvex (combinatorial) models. A general quasi-equivalence between these problems is established in the sense explained next. There exists a strictly decreasing sequence of critical values $\{\beta_k\}$ that partitions the positive axis into a certain number of intervals. For every β inside the K-th interval, the regularized problem and the K-constrained problem share exactly the same set of optimal solutions. For $\beta = \beta_k$, the optimal set of the regularized problem is composed of the optimal solutions of the K-constrained problem and the constrained problem corresponding to the next interval (and possibly a few others in-between). All β_k's are obtained from the optimal values of the constrained problems. The strict decay of $\{\beta_k\}$ is a necessary and sufficient condition. We will present all important points concerning this quasi-equivalence.

Small-size exact numerical tests illustrate the theoretical results. By way of conclusion, the K-sparsity problem offers wider possibilities which is not necessarily an advantage.

A2 - December 11, 17:00 – 17:25

Simultaneous high dynamic range image reconstruction and denoising for non-static scenes

Pablo Musé

Universidad de la República, Uruguay

pmuse@fing.edu.uy

The human eye has the ability to capture scenes of very high dynamic range, retaining details in both dark and bright regions. This is not the case for current standard digital cameras. The limited capacity of the sensor cells makes it impossible to record the irradiance from very bright regions for long exposures (saturated pixels) while very few photons will be captured in the dark regions for short exposures (noisy pixels).

Since the seminal work of Mann and Picard in 1994, the standard way to build high dynamic range (HDR) images from regular cameras has been to combine a reduced number of photographs captured with different exposure times. The algorithms proposed in the literature differ in the strategy used to combine these frames.

Under the hyphotesis of perfectly aligned images (fixed scene and static camera), a study of the different noise sources in the image acquisition process allows us to model the image fusion as a statistical estimation problem. We derive theoretical bounds for the performance of the HDR estimation problem and show that, even with a small number of photographs, the maximum likelihood estimator performs extremely close to these bounds.

In practice, scenes are dynamic and images are usually acquired with a hand-held camera. We propose a new HDR image generation approach that simultaneously copes with these problems and exploits image redundancy to produce a denoised result. A reference image is chosen and a patch-based approach is used to find similar pixels that are then combined for the irradiance estimation. This patch-based approach permits to obtain a denoised result and is robust to image misalignments and object motion. Results

show significant improvements in terms of noise reduction over previous HDR image generation techniques, while being robust to motion and changes between the exposures.

Joint work with Cecilia Aguerrebere (Duke University, USA), Julie Delon (Télécom ParisTech, France) and Yann Gousseau (Télécom ParisTech, France).

A2 - December 11, 17:30 – 17:55

Denoising an Image by Denoising its Curvature

Stacey Levine
Duquesne University, USA
sel@mathcs.duq.edu

In this work we argue that when an image is corrupted by additive noise, its curvature image is less affected by it. In particular, we demonstrate that for sufficient noise levels, the PSNR of the curvature image is larger than that of the original image. This leads to the speculation that given a denoising method, we may obtain better results by applying it to the curvature image and then reconstructing from it a clean image, rather than denoising the original image directly. Numerical experiments confirm this for several PDE-based and patch-based denoising algorithms.

Joint work with Marcelo Bertalmío (Universitat Pompeu Fabra).

A2 - December 11, 18:00 – 18:25

Recent algorithmic and theoretical advances on graph matching.

Marcelo Fiori
Universidad de la República, Uruguay
mfiori@fing.edu.uy

Problems related with graph matching are very important both from a theoretical and practical point of view, with several applications from image and video analysis to biological and biomedical problems. In this talk we will cover both aspects of this challenging problem.

First, a graph matching algorithm will be presented, combining sparsity ideas with relaxations of the graph matching problem, yielding a robust method, particularly well suited for multimodal data. Then, new probabilistic results about convex and non-convex relaxations of the graph matching problem will be discussed, also suggesting some practical considerations. Finally, new spectral results will be presented, proving the equivalence for the graph matching problem and a relaxed version for certain graphs, also shedding light to the relationship between spectral properties and the automorpshism group of a graph.

Joint work with Guillermo Sapiro (Duke University, USA), Pablo Sprechmann (New York University, USA), Joshua Vogelstein (Duke University, USA), Pablo Musé (Universidad de la República, Uruguay), Vince Lyzinski (Johns Hopkins University, USA), Donniell Fishkind (Johns Hopkins University, USA) and Carey E. Priebe (Johns Hopkins University, USA)..

A2 - December 11, 18:30 – 18:55

Matrix recovery from coarse observations

Mark Davenport
Georgia Institute of Technology, USA
mdav@gatech.edu

In this talk I will describe a theory and techniques for solving matrix completion and similar problems when given extremely coarse (e.g., 1-bit) observations. As an example, instead of observing a subset of the real-valued entries of a matrix M, we might obtain a small number of binary (1-bit) measurements generated according to a probability distribution determined by the real-valued entries of M. The central question I will discuss is whether or not it is possible to obtain an accurate estimates of M from data of this form. In general this would seem impossible, but I will show that the maximum likelihood estimate under a suitable constraint returns an accurate estimate of M under certain natural conditions. I will conclude by discussing several extensions and applications of these techniques to similar problems.

Joint work with Yaniv Plan (University of British Columbia), Ewout van den Berg (IBM Watson) and Mary Wootters (Carnegie Mellon University).

A2 - December 12, 14:30 – 14:55

FUNDAMENTALS OF DYNAMICAL SAMPLING

Carlos Cabrelli
Dept. of Mathematics, FCEyN, UBA and IMAS-UBA-CONICET, Argentina
cabrelli@dm.uba.ar

Dynamical sampling refers to the process that results from sampling an evolving signal f at various times. The fundamental question of this spatial-temporal sampling is: when do coarse samplings taken at varying times contain the same information as a finer sampling taken at the earliest time? In other words, under what conditions on an evolving system, can time samples be traded for spatial samples?

Because dynamical sampling uses samples from varying time levels for a single reconstruction, it departs from classical sampling theory in which a signal f does not evolve in time and is to be reconstructed from its samples at a single time $t = 0$.

In this talk we study this problem in finite dimensional spaces, and for a large class of self adjoint operators in infinite dimensional spaces.

Joint work with Akram Aldroubi (Vanderbilt University, US), Ursula Molter (University of Buenos Aires, Argentina), Sui Tang (Vanderbilt University, US)..

A2 - December 12, 15:00 – 15:25

ON SPECTROGRAM LOCAL MAXIMA

Patrick Flandrin
CNRS & Ecole normale supérieure de Lyon, France
flandrin@ens-lyon.fr

In close connection with time-frequency uncertainty relations, spectrograms have some built-in redundancy which constrains the landscape of their surface, thus calling for simplified descriptions based on a reduced number of salient features. This is investigated in some detail for the distribution of local maxima in the generic case of white Gaussian noise. A simple model, based on a randomized hexagonal lattice

structure, is proposed for such a distribution considered as a spatial point process in the time-frequency plane. The rationale of the model is discussed, its relevance is tested with respect to the cumulative distribution function of nearest-neighbour distance between local maxima, and the deviation from the complete spatial randomness of an equivalent Poisson model is quantified. Attaching a Voronoi tessellation to local maxima ends up with a constrained distribution of cells that reflects uncertainty and paves the way for the modeling of spectrogram enhancements such as those offered by reassignment or synchrosqueezing.

A2 - December 12, 15:30 – 15:55

High dimensional learning rather than computing in quantum chemistry

Matthew Hirn
Ecole normale superieure, France
matthew.hirn@ens.fr

Physical functionals are usually computed as solutions of variational problems or from solutions of partial differential equations, which may require huge computations for complex systems. Quantum chemistry calculations of molecular ground state energies is such an example. Machine learning algorithms do not simulate the physical system but estimate solutions by interpolating values provided by a training set of known examples. However, precise interpolations may require a number of examples that is exponential in the system dimension, and are thus intractable. This curse of dimensionality may be avoided by computing interpolations in smaller approximation spaces, which take advantage of physical invariants. We introduce deep multiscale learning architectures in a similar vein to deep neural networks, which compute such invariant approximations via iterated wavelet transforms. Theoretical results relating these architectures to the Coulomb potential from classical physics will motivate numerical applications for molecular energies in quantum chemistry, in relation with Density Functional Theory.

Joint work with Stephane Mallat (Ecole normale superieure, France) and Nicolas Poilvert (Pennsylvania State University, USA).

A2 - December 12, 16:00 – 16:25

Estimation of bandlimited stochastic operators using SIC-POVMs

Götz Pfander
Jacobs University Bremen / Katholische Universität Eichstätt, Germany
g.pfander@jacobs-university.de

We describe a novel estimator for bandlimited stochastic operators. The operators considered are frequently used to model randomly varying parts of radar targets and communications channels. The presented estimator is based on insights on stochastic modulation spaces, in operator sampling theory, and on finite dimensional Gabor frames. In case of wide sense stationary with uniform scattering (WSSUS) targets, we exhibit a connection to symmetrically information complete positive operator valued measures (SIC-POVMs) as considered in quantum information theory.

Joint work with Pavel Zheltov (Jacobs University Bremen).

A2 - December 12, 17:00 – 17:25

Consistency of probability measure quantization by means of power repulsion–attraction potentials

Massimo Fornasier

Technische Universität München, Germany

massimo.fornasier@ma.tum.de

In this talk we are concerned with the study of the consistency of a variational method for probability measure quantization, deterministically realized by means of a minimizing principle, balancing power repulsion and attraction potentials. The proof of consistency is based on the construction of a target energy functional whose unique minimizer is actually the given probability measure to be quantized. Then we show that the discrete functionals, defining the discrete quantizers as their minimizers, actually Gamma-converge to the target energy with respect to the narrow topology on the space of probability measures. A key ingredient is the reformulation of the target functional by means of a Fourier representation, which extends the characterization of conditionally positive semi-definite functions from points in generic position to probability measures.

Joint work with Jan-Christian Hütter (MIT, USA).

A2 - December 12, 18:05 – 18:55

Weighted sparsity for function approximation and interpolation

Rachel Ward

University of Texas at Austin, USA

rward@math.utexas.edu

Functions of interest are often smooth and sparse in some sense, and both priors should be taken into account when interpolating sampled data. Linear interpolation methods are effective under strong regularity assumptions, but do not incorporate nonlinear sparsity structure. At the same time, nonlinear methods such as ℓ_1 minimization can reconstruct sparse functions from very few samples, but do not encourage smoothness without adding weights. We argue that weighted ℓ_1 minimization effectively merges the two approaches, promoting both sparsity and smoothness in reconstruction. Along the way, we introduce a notion of weighted sparsity and extend concepts from compressive sensing such as the restricted isometry property and null space property to accommodate weighted sparse expansions. We expect these developments to be of independent interest in the study of structured sparse approximations and continuous-time compressive sensing problems.

Joint work with Holger Rauhut (Aachen University).

A2 - December 13, 14:30 – 14:55

Color Stabilization Along Time and Across Shots of the Same Scene, for One or Several Cameras of Unknown Specifications

Marcelo Bertalmío

Universitat Pompeu Fabra, Spain

marcelo.bertalmio@upf.edu

We propose a method for color stabilization of shots of the same scene, taken under the same illumination, where one image is chosen as reference and one or several other images are modified so that their colors match those of the reference. We make use of two crucial but often overlooked observations: firstly, that the core of the color correction chain in a digital camera is simply a multiplication by a 3x3 matrix; secondly, that to color-match a source image to a reference image we don't need to compute their two color correction matrices, it's enough to compute the operation that transforms one matrix into the other. This operation is a 3x3 matrix as well, which we call H. Once we have H, we just multiply by it each pixel value of the source and obtain an image which matches in color the reference. To compute H we only require a set of pixel correspondences, we don't need any information about the cameras used, neither models nor specifications or parameter values. We propose an implementation of our framework which is very simple and fast, and show how it can be successfully employed in a number of situations, comparing favourably with the state of the art. There is a wide range of applications of our technique, both for amateur and professional photography and video: color matching for multi-camera TV broadcasts, color matching for 3D cinema, color stabilization for amateur video, etc.

Joint work with Javier Vazquez-Corral (Universitat Pompeu Fabra, Spain).

A2 - December 13, 15:00 – 15:25

Multi-level structured sparse models

Pablo Sprechmann
Courant Institute, New York University, United States
pablo.sprechmann@nyu.edu

Parsimony, including sparsity and low rank, has been shown to successfully model data in numerous machine learning and signal processing tasks. Sparse models assume minimal prior knowledge about the data, asserting that the signal has many coefficients close or equal to zero when represented in a given domain. From a data modeling point of view, sparsity can be seen as a form of regularization, that is, as a device to restrict or control the set of coefficient values which are allowed in the model to produce an estimate of the data. While this model has been proven to be effective in many settings, it is found insufficient in difficult inverse problems such as source separation. In these situations, one could greatly benefit from learning further structure present in the data. We propose a new type of structure sparse models that aim at learning a multi-level sparse representation of the data. The proposed representation is organized hierarchically and aims at learning high-level structure, such as dependencies (or correlations) in the activations and short-term temporal dynamics. We evaluate the proposed model on a monaural audio separation task and discuss connections with deep learning.

Joint work with Joan Bruna (New York University, USA) and Yann Lecun (New York University, USA).

A2 - December 13, 15:35 – 16:25

On the stability of least-squares approximations. Application in acoustics.

Albert Cohen
Université Paris VI, France
cohen@ann.jussieu.fr

A2 - December 13, 17:00 – 17:25

Texture aware video inpainting of complex scenes

Andrés Almansa
CNRS & Telecom ParisTech, France
andres.almansa@telecom-paristech.fr

We present an automatic video inpainting algorithm which relies on the multi-scale optimisation of a patch-based functional enforcing self-similarity. Unlike previous approaches, the best patch candidates are selected using persistent multi-scale texture attributes. We show that this rationale prevents the usual wash-out of textured and cluttered parts of videos. The resulting approach is able to successfully and automatically inpaint complex situations, including high resolution sequences with dynamic textures and multiple moving objects.

Joint work with Alasdair Newson (Duke University), Matthieu Fradet (Technicolor), Yann Gousseau (Telecom ParisTech) and Patrick Perez (Technicolor).

A2 - December 13, 17:30 – 17:55

Self similarity and spectral correlation adaptive algorithm for image interpolation

Antoni Buades
Universitat Illes Balears, Spain
toni.buades@uib.es

Most common cameras use a CCD sensor device measuring a single color per pixel. The other two color values of each pixel must be interpolated from the neighboring pixels in the so-called demosaicking process. Recent devices include also grey level sensors in order to increase the signal to noise ratio. These recent devices get closer to the satellite acquisition systems, where the color and luminance information are acquired separately, and with different resolutions; a high-resolution luminance image and a low-resolution multispectral one.

In this talk, I will review classical solutions to these two different interpolation problems and present new algorithms.

Joint work with J. Duran (Universitat Illes Balears, Spain).

A2 - December 13, 18:00 – 18:25

A one stage Sigma-Delta decoder for compressed sensing measurements

Rongrong Wang
University of British Columbia, Canada
rongwang@math.ubc.ca

We analyze how efficiently Sigma-Delta quantization works for quantizing compressed (sub-Gaussian) measurements of sparse and compressible signals. To this end, we propose a one-stage reconstruction algorithm based on convex optimization that yields consistent reconstruction. The algorithm works in the cases of fine and coarse quantization including one-bit quantization, with a reconstruction error decaying inverse polynomially in the quantization order. We show that this decay rate is nearly optimal

among all possible reconstruction algorithms by a geometric argument about quantization cells. When we optimize over all quantization orders, the algorithm can achieve root exponential error decay with respect to the "oversampling factor". Finally, we show that by further compressing the quantized data via a Johnson-Lindenstrauss embedding, exponential decay (as a function of the total bit budget) is achieved.

Joint work with This is joint work with Rayan Saab (UCSD) and Ozgur Yilmaz (UBC).

A2 - December 13, 18:30 – 18:55

Orthonormal Bases Generated by Cuntz Algebras

Myung-Sin Song
Southern Illinois University Edwardsville, USA
msong@siue.edu

We show how some orthonormal bases can be generated by representations of the Cuntz algebra such as Fourier bases on fractal measures, generalized Walsh bases on the unit interval and piecewise exponential bases on the middle third Cantor set.

Joint work with Dorin Dutkay (University of Central Florida, USA) and Gabriel Picioroaga (University of South Dakota, USA).

A2 - Poster

Local convergence of an algorithm for subspace identification with missing data

Laura Balzano
University of Michigan, USA
girasole@umich.edu

Low-dimensional linear subspace approximations to high-dimensional data find application in a great variety of applications where missing data are the norm, not only because of errors and failures in data collection, but because it may be impossible to collect and process all the desired measurements.

In this poster, I will describe recent results on estimating subspace projections from incomplete data. I will discuss the convergence guarantees and performance of the algorithm GROUSE (Grassmannian Rank-One Update Subspace Estimation), a subspace tracking algorithm that performs gradient descent on the Grassmannian. I will also discuss the relationship of GROUSE with an incremental SVD algorithm, and show results of GROUSE applied to problems in computer vision.

Joint work with Stephen J. Wright.

A2 - Poster

Finitely generated shift invariant spaces with extra invariance nearest to observed data

Carolina Mosquera
Universidad de Buenos Aires, CONICET, Argentina
mosquera@dm.uba.ar

Let $m, \ell \in \mathbb{N}$, M be a closed subgroup of $\mathbb{R}^d$ containing $\mathbb{Z}^d$ and $F = \{f_1, \ldots, f_m\} \subset L^2(\mathbb{R}^d)$. We study the problem of finding the shift invariant space V of length less or equal to ℓ which is also $M-$invariant such that V is "closest" to the functions F in the sense that

$$V = argmin_{V' \in V_M^\ell} \sum_{j=1}^m \|f_j - P_{V'} f_j\|^2,$$

where V_M^ℓ is the set of all shift invariant spaces V' of length less or equal to ℓ which are also $M-$invariant, and $P_{V'}$ is the orthogonal projection on V'.

Also we consider this problem for a particular set of translation invariant spaces.

Joint work with Carlos Cabrelli (Universidad de Buenos Aires, CONICET, Argentina).

A2 - Poster

Relax, no need to round: Integrality of clustering formulations.

Soledad Villar
University of Texas at Austin, United States
mvillar@math.utexas.edu

When dealing with combinatoric NP-hard problems one can either find an efficient algorithm that works in every case but its solution is not optimal (this is the philosophy behind approximation algorithms) or efficiently compute a probably optimal solution of the original problem under more restrictive hypothesis (which is the spirit of compressed sensing recovery guarantees).

In this work we take the second approach applied to point cloud clustering problems; showing recovery conditions of an integral solution through their convex relaxations. We focus on two of the most common optimization problems for unsupervised clustering: k-means and k-medians clustering.

Joint work with P. Awasthi (Princeton University), A. S. Bandeira (Princeton University), M. Charikar (Princeton University), R. Krishnaswamy (Princeton University) and R. Ward (University of Texas at Austin).