ITWIST: INTERNATIONAL TRAVELING WORKSHOP ON INTERACTIONS BETWEEN LOW-COMPLEXITY DATA MODELS AND SENSING TECHNIQUES

CIRM, LUMINY, 19 - 23 NOVEMBER, 2018

Abstracts

Ben Adcock (Simon Fraser University, Canada)

Compressed sensing and high-dimensional approximation: theory and applications

Abstract: Many problems in computational science require the approximation of a high-dimensional function from limited amounts of data. For instance, a common task in Uncertainty Quantification (UQ) involves building a surrogate model for a parametrized computational model. Complex physical systems involve computational models with many parameters, resulting in multivariate functions of many variables. Although the amount of data may be large, the curse of dimensionality essentially prohibits collecting or processing enough data to reconstruct such a function using classical approximation techniques. Over the last five years, spurred by its successful application in signal and image processing, compressed sensing has begun to emerge as potential tool for surrogate model construction UQ. In this talk, I will give an overview of application of compressed sensing to high-dimensional approximation. I will demonstrate how the appropriate implementation of compressed sensing overcomes the curse of dimensionality (up to a log factor). This is based on weighted l1 regularizers, and structured sparsity in so-called lower sets. If time, I will also discuss several variations and extensions relevant to UQ applications, many of which have links to the standard compressed sensing theory. These include dealing with corrupted data, the effect of model error, functions defined on irregular domains and incorporating additional information such as gradient data. I will also highlight several challenges and open problems.

Laurent Daudet (Institut Langevin, Paris)

Optical random features for large-scale machine learning

Abstract: The propagation of coherent light through a thick layer of scattering material is an extremely complex physical process. However, it remains linear, and under certain conditions, if the incoming beam is spatially modulated to encode some data, the output as measured on a sensor can be modeled as a random projection of the input, i.e. its multiplication by an iid random matrix. One can leverage this principle for compressive imaging, and more generally for any data processing pipeline involving large-scale random projections. This talk will present a series of proof of concept experiments in machine learning, and discuss recent technological developments of optical co-processors within the startup LightOn.

Mike Davies (University of Edinburgh, UK)

Inexact Gradient Projection and Fast Data Driven Compressed Sensing: theory and application

Abstract: We consider the convergence of the iterative projected gradient (IPG) algorithm for arbitrary (typically nonconvex) sets and when both the gradient and projection oracles are only computed approximately. We consider different notions of approximation of which we show that the Progressive Fixed Precision (PFP) and (1+epsilon) optimal oracles can achieve the same accuracy as for the exact IPG algorithm. We also show that the former scheme is also able to maintain the (linear) rate of convergence of the exact algorithm, under the same embedding assumption, while the latter requires a stronger embedding condition, moderate compression ratios and typically exhibits slower convergence. We apply our results to accelerate solving a class of data driven compressed sensing problems, where we replace iterative exhaustive searches over large datasets by fast approximate nearest neighbour search strategies based on the cover tree data structure. Finally, if there is time we will give examples of this theory applied in practice for rapid enhanced solutions to an emerging MRI protocol called magnetic resonance fingerprinting for quantitative MRI.

Sjoerd Dirksen (RWTH Aachen, Germany)
Robust one-bit compressed sensing with non-Gaussian measurements

Abstract: In the traditional compressed sensing literature, it is implicitly assumed that one has direct access to noisy analog linear measurements of an (unknown) signal. In reality, these analog measurements need to be quantized to a finite number of bits before they can be transmitted, stored, and processed. In the emerging theory of quantized compressed sensing it is studied how to jointly design a quantizer, measurement procedure, and reconstruction algorithm in order to accurately recover low-complexity signals. In the popular one-bit compressed sensing model, each linear analog measurement is quantized to a single bit in a memoryless fashion. This quantization operation can be implemented with energy-efficient hardware. There is by now a rich theory available for one-bit compressed sensing with standard Gaussian measurements. Outside of this purely Gaussian setting, very little is known about one-bit compressed sensing. In fact, recovery can in general easily fail for non-Gaussian measurement matrices, even if they are known to perform optimally in "unquantized" compressed sensing. In my talk, I will show that this picture completely changes if we use dithering, i.e., deliberately add noise to the measurements before quantizing them. By using well-designed dithering, it becomes possible to accurately reconstruct low-complexity signals from a small number of one-bit quantized measurements, even if the measurement vectors are drawn from a heavy-tailed distribution. The reconstruction results that I will present are very robust to noise on the analog measurements as well as to adversarial bit corruptions occurring in the quantization process. If the measurement matrix is subgaussian, then accurate recovery can be achieved via a convex program. The proofs of these reconstruction theorems are based on novel random hyperplane tessellation results.

Based on joint work with Shahar Mendelson (Technion Haifa/ANU Canberra).

Simon Foucart (Texas A&M University, USA)
Standard, One-Bit, and Saturated Compressive Sensing

Abstract: In this talk, I shall summarize recent results about sparse recovery from compressive, possibly nonlinear, measurements. The focus is put on ℓ_1-minimization and (iterative) hard thresholding as recovery procedures and a modification of the classic restricted isometry property appears as a common tool throughout. In the standard scenario, the measurements made on a sparse vector $\mathbf{x} \in \mathbb{R}^N$ take the form $\mathbf{y} = \mathbf{A}\mathbf{x} \in \mathbb{R}^m$ for some matrix $\mathbf{A} \in \mathbb{R}^N$ with $m \ll N$. Exact recovery of $\mathbf{x}$ is achievable in this case. In the one-bit scenario, the measurements are quantized to the extreme as $\mathbf{y} = \mathrm{sgn}(\mathbf{A}\mathbf{x}) \in \{\pm 1\}^m$. Only approximate recovery of the direction of $\mathbf{x}$ is possible in this case and the goal is to quantify the recovery error. In the saturated scenario, the measurements take the form $\mathbf{y} = \mathcal{S}(\mathbf{A}\mathbf{x}) \in [-\mu, \mu]^m$ for some saturation function $\mathcal{S}$ with parameter $\mu > 0$. As intuitively expected, the results bridge the standard theory ($\mu \to \infty$) and the one-bit theory ($\mu \to 0$). Indeed, as μ decreases, there is a regime where exact recovery is achievable, followed by a regime where approximate recovery of $\mathbf{x}$ is possible, which transitions into a regime where the one-bit situation prevails.

Alexandre Gramfort (Inria, UniversitéParis-Saclay, France)

Optimization strategies for fast inverse problems under sparsity constraints (with some applications in neuroimaging)

Abstract: In this course AG will:

- Motivate the use of sparse regularizations in the context of neuroscience applications;
- Review some results on coordinate descent methods starting from more well known iterative algorithms such as a proximal or projected gradient descent;
- Cover the dual construction of Lasso-type solvers and explain how it can be used to control optimality, derive accelerations with screening rules and working set methods;
- Present how such methods can be efficiently implemented in Python using Numba or Cython (as done in the Scikit-Learn software).

Laurent Jacques (UCLouvain, Belgium)

Quantized compressed sensing and related data embeddings

Abstract: In this course we will discover the interplay of Compressive Sensing theory, as introduced by Simon Foucart in this doctoral school, with the unavoidable quantization of any sensing procedure, that is, the standard analog-to-digital conversion operated in actual sensing devices in order to efficiently transmit, store or process recorded data. This interaction will lead us to the definition of interesting mathematical questions in high dimensional geometry, with for instance the study of certain embedding properties for (1-bit) quantized random projections, i.e., the preservation of pairwise vector distance in the quantized and projected domain, up to controllable distortions. In particular, this course will cover the following aspects:

- Early attemps to combine CS and quantization;
- Principles of memoryless scalar quantization: 1-bit and multi-bits;
- Consistent reconstruction methods, in quantization theory and in quantized compressive sensing;
- 1-bit Compressive Sensing: compatible sensing matrices, reconstruction algorithms and guarantees;
- Multi-bit Quantized Compressive Sensing and embeddings: the benefit of dithering;
- Overview of other quantization methods, e.g., noise shaping quantization and Sigma-Delta QCS.

Ulugbek Kamilov (Washington University, USA)

Signal Processing for Nonlinear Diffractive Imaging: Acquisition, Reconstruction, and Applications

Abstract: Can modern signal processing be used to overcome the diffraction limit? The classical diffraction limit states that the resolution of a linear imaging system is fundamentally limited by one half of the wavelength of light. This implies that conventional light microscopes cannot distinguish two objects placed within a distance closer than $0.5 \times 400 = 200nm$ (blue) or $0.5 \times 700 = 350nm$ (red). This significantly impedes biomedical discovery by restricting our ability to observe biological structure and processes smaller than $100nm$. Recent progress in sparsity-driven signal processing has created a powerful paradigm for increasing both the resolution and overall quality of imaging by promoting model-based image acquisition and reconstruction. This has led to multiple influential results demonstrating super-resolution in practical imaging systems. To date, however, the vast majority of work in signal processing has neglected the fundamental nonlinearity of the object-light interaction and its potential to lead to resolution enhancement. As a result, modern theory heavily focuses on linear measurement models that are truly effective only when object-light interactions are weak. Without a solid signal processing foundation for understanding such nonlinear interactions, we undervalue their impact on information transfer in the image formation. This ultimately limits our capability to image a large class of objects, such as biological tissue, that generally are in large-volumes and interact strongly and nonlinearly with light.

The goal of this talk is to present the recent progress in model-based imaging under multiple scattering. We will discuss several key applications including optical diffraction tomography, Fourier Ptychography, and large-scale Holographic microscopy. We will show that all these application can benefit from models, such as the Rytov approximation and beam propagation method, that take light scattering into account. We will discuss the integration of such models into the state-of-the-art optimization algorithms such as FISTA and ADMM. Finally, we will describe the most recent work that uses learned-priors for improving the quality of image reconstruction under multiple scattering.

Gitta Kutyniok (TU Berlin, Germany)

Compressed Sensing from an Analysis Viewpoint: Successes and Failures

Abstract: Compressed Sensing is based on the assumption that the signal of interest exhibits some low-complexity behavior in the sense of sparsity. Since typically a transform is necessary to reveal this property, the analysis sparsity (cosparsity) model has gained increasing attention. To incorporate this into the recovery algorithm, an ℓ_1-analysis-minimization strategy is often pursued, which is very successful, for instance, for multiscale transforms. But despite this empirical success, many theoretical properties of the analysis approach remained unexplored.

In the first part of this lecture, we will discuss a generalized notion of sparsity, which allows us to derive very precise recovery guarantees for ℓ_1-analysis-minimization, enabling accurate predictions of its sample complexity. The corresponding bounds on the number of required measurements do explicitly depend on the Gram matrix of the analysis operator and therefore particularly account for its mutual coherence structure. Our findings surprisingly defy conventional wisdom which promotes the sparsity of analysis coefficients as the crucial quantity to study.

However, in certain applications due to a lack of data such strategies fail as in limited-angle computed tomography. For such situations we will present a general strategy to combine ℓ_1-analysis-minimization with deep learning. In this approach, learning is only targeted to those parts, which ℓ_1-analysis-minimization is incapable to handle, thereby still allowing a maximal control on the recovery procedure. We will also present numerical experiments showing the superiority of such a strategy.

Denali Molitor (UCLA, USA)

A simple approach to hierarchical classification

Abstract: We extend a recent simple classification approach for binary data in order to efficiently classify hierarchical data. In certain settings, specifically, when some classes are significantly easier to identify than others, we showcase computational and accuracy advantages.