



Off-the-grid methods for Inverse Problems in Imaging

October 21–22 2023

Institut Henri Poincaré (Amphithéâtre Hermite), Paris

Summary

Variational methods are a flexible and powerful tool for the resolution of inverse problems in imaging. Contrary to a traditional approach which consists in reconstructing an image on a pixel grid, yielding a discrete problem, several recent methods such as the Beurling LASSO are directly formulated in a continuous domain. This generally yields solutions which are more faithful to the underlying physical problem, and the resulting discretizations may exploit the structures of the solutions (sparsity) and the continuous (or even differentiable) nature of the problem.

The aim of this workshop is to provide an overview of the latest advances on continuous methods for imaging.

Program

Tuesday November 21

09h30 – 10h00	<i>Welcoming coffee</i>	
10h00 – 10h20	<i>Opening</i>	
10h20 – 11h10	Kristian Bredies	Models and algorithms for off-the-grid point tracking in dynamic inverse problems
11h10 – 12h00	Bastien Laville	Gridless curve reconstruction: divergence regularisation and untangling by (sub)Riemannian metric

12h00 – 13h30	<i>Lunch break</i>	
----------------------	--------------------	--

13h30 – 14h20	Yann Traonmilin	A few years of non-convex off-the-grid estimation
14h20 – 15h10	Lénaïc Chizat	Global vs Local for Optimization in the Space of Measure

15h10 – 15h40	<i>Coffee break</i>	
----------------------	---------------------	--

15h40 – 16h30	Clément Hardy	Prediction and testing of mixtures of features issued from a continuous dictionary
16h30 – 17h20	Clément Elvira	Safe screening for total variation norm

Wednesday November 22

09h00 – 09h50	Maxime Ferreira Da Costa	Extremal approximations in the bandlimit and the Rayleigh criterion for super-resolution
9h50 – 10h20	<i>Coffee break</i>	
10h20 – 11h10	Marta Lazzaretti	Off-the-grid regularisation for Poisson inverse problems
11h10 – 12h00	Daniel Walter	Towards optimal sensor placement for inverse problems in spaces of measures

12h00 – 13h30	<i>Lunch break</i>	
----------------------	--------------------	--

13h30 – 14h20	Clarice Poon	Super-resolved Lasso
14h20 – 15h10	Karl Hajjar	TBA

15h10 – 15h40	<i>Coffee break</i>	
----------------------	---------------------	--

15h40 – 16h30	Blanche Buet	Flagfolds: multi-dimensional varifolds to handle discrete surfaces
16h30 – 17h20	Romain Petit	Exact recovery of the support of piecewise constant images via total variation regularization

Organizing committee

Claire Boyer (Sorbonne Université)
Jean-Baptiste Courbot (Université de Haute Alsace)
Yohann De Castro (Institut Camille Jordan)
Vincent Duval (INRIA)
Emmanuel Soubies (CNRS, Université de Toulouse)

Funding

ANR CIPRESSI

Abstracts

Models and algorithms for off-the-grid point tracking in dynamic inverse problems

Kristian Bredies (*University of Graz*)

We present and discuss strategies for the solution of dynamic inverse problems in spaces of measures where for each time point, a time-dependent linear forward operator mapping measures to time-dependent Hilbert-space data has to be inverted. These problems are regularized with dynamic optimal-transport energies that base on the continuity equation as well as convex functionals of Benamou—Brenier and Hellinger–Kantorovich type [ESAIM:M2AN 54(6):2351—2382, 2020]. Sparsity results then give rise to the study of the extremal points of the Benamou—Brenier/Hellinger–Kantorovich energy subject to the continuity equation. For the latter, it is proven that the extremal points are realized by point masses moving along curves with Sobolev regularity [Bull. LMS 53(5):1436—1452, 2021] [CPDE 47(10):2023–2069, 2022]. This result will be employed to develop numerical optimization algorithms of generalized conditional gradient type for off-the-grid point tracking. We present instances of this algorithm that are tailored towards this task in the context of dynamic inverse problems. Finally, the application and numerical performance of the method is demonstrated for sparse dynamic superresolution [FOCM 23(3):833–898, 2023].

Gridless curve reconstruction: divergence regularisation and untangling by (sub)Riemannian metric

Bastien Laville (*Université Côte d’Azur*)

Recent years have seen the development of super-resolution variational optimisation in measure spaces. These so-called off-the-grid approaches offer both theoretical and numerical results, with very convincing results in biomedical imaging. However, the gridless variational optimisation is generally formulated for reconstruction of point sources, which is not always suitable for biomedical imaging applications: more realistic biological structures such as curves should also be reconstructed. In the first part of this talk, we propose a new strategy for the reconstruction of curves in an image through an off-the-grid variational framework, thanks to the sharp characterisation of the extreme points of the unit ball of a new regulariser thus enabling new theoretical and numerical results for optical imaging. In a second part of the talk, we investigate a new strategy for off-the-grid curve untangling. Indeed, some latest developments in the gridless community enabled the reconstruction of moving Dirac, understood as point sources following some paths along the time. It was nevertheless by design difficult to recover crossing curves, a situation yet arising quite often in practical biomedical problem. We then present a new strategy, using a lift on the roto-translational space and a sub-Riemannian metric regularisation to untangle this crossing paths, with some practical results for Localisation Microscopy.

A few years of non-convex off-the-grid estimation

Yann Traonmilin (*CNRS & Institut de Mathématiques de Bordeaux*)

In this talk, we focus on non-convex approaches for off-the-grid spike estimation. Centered around the study of basins of attraction of a non-convex functional, we explain how the study of recovery guarantees can be generally linked with the number of available measurements. With a general result on non-convex estimation of low-dimensional models, we show that the size of basins of attraction explicitly increases with respect to the number of measurements, with tight bounds for spikes recovery. These results lead to the conception of a fast algorithm for the recovery of many off-the-grid spikes: over-parametrized projected gradient descent (OP-PGD), showing promising results on realistic datasets. We also are able to give a theoretical partial control of the quality of continuous orthogonal matching pursuit without sliding which is the initialization procedure of OP-PGD.

Global vs Local for Optimization in the Space of Measure

Lénaïc Chizat (*EPFL*)

The difficulty of optimization in the space of measures on a continuous domain can be separated into two problems, each corresponding to a phase of optimization algorithms: (i) the global phase : how to approach the neighborhood of a global minimizer efficiently? and (ii) the local phase : how to obtain high-accuracy solutions from a warm start? For the global phase, we'll discuss approaches based on convex optimization – presenting in particular an accelerated mirror descent method that converges provably faster than FISTA or Frank-Wolfe – and approaches based on non-linear Langevin dynamics, with new convergence guarantees. For the local phase, we will argue, via quantitative convergence guaranties, that following the Wasserstein-Fisher-Rao geometry is the best way to “slide” the grid/particles with first order methods.

Prediction and testing of mixtures of features issued from a continuous dictionary

Clément Hardy (*École des Ponts ParisTech*)

In this talk, we will consider observations that are random elements of a Hilbert space resulting from the sum of a deterministic signal and a noise. The signals considered will be linear combinations (or mixtures) of a finite number of features issued from continuous parametric dictionaries.

In order to estimate the linear coefficients as well as the non-linear parameters of a mixture in the presence of noise, we propose estimators that are solutions to an optimization problem. We shall quantify the performance of these estimators with respect to the quality of the observations by establishing prediction and estimation bounds that stand with high probability. In practice, it is common to have a set of observations (possibly a continuum) sharing common features. The question arises whether the estimation of signals can be improved by taking advantage of their common structure. We give a framework in which this improvement occurs.

Next, we shall test whether a noisy observation is derived from a given signal and give non-asymptotic upper bounds for the associated testing risk. In particular, our test encompasses the signal detection framework. We will derive an upper bound for the strength that a signal must have in order to be detected in the presence of noise.

Safe screening for total variation norm

Clément Elvira (*CentraleSupélec Rennes*)

Total variation regularized optimization problems over the space of Radon measures have gained a lot of interest in the last decade. One element underlying this success is the ‘discrete’ property of their solutions, which consist of a finite number of (weighted) Dirac masses. However, a bottleneck of standard algorithmic solutions is the need for extensive (and iterative) exploration of the parameter space. In this study, we show that the rationale behind “safe screening” –a celebrated acceleration technique for sparsity promoting optimization problems– can be extended to the context of radon measures. More precisely, we describe practical procedures that allow to test whether any solution contains Dirac masses that lie in a given subset of the parameter space. If the test is passed, one can thus safely discard the corresponding subset of parameters, that is, without affecting the solutions. Finally, we show on numerical examples that our procedures allow to accelerate several existing algorithms.

Joint work with Thu-Le Tranh, Hong-Phuong Dang and Cédric Herzet.

Extremal approximations in the bandlimit and the Rayleigh criterion for super-resolution

Maxime Ferreira Da Costa (*CentraleSupélec*)

Of particular interest in imaging sciences and telecommunications is the super-resolution problem, which consists in recovering a stream of spikes (point sources) from the noisy observation of a few number of its first trigonometric moments weighted by the ones of the point-spread function (PSF). The empirical feasibility of this problem has been known since the work of Rayleigh on diffraction to be essentially driven by the separation between the spikes to recover.

We present a novel statistical framework based on the spectrum of the Fisher information matrix (FIM) to quantify the stability limit of super-resolution as a function of the PSF. In the regime where the minimal separation is inversely proportional to the number of acquired moments, we show the existence of a separation constant above which the minimal eigenvalue of the FIM is not asymptotically vanishing—defining a statistical resolution limit. Notably, a relationship between the total variation of the autocorrelation function of the PSF and its association resolution limit is highlighted.

Off-the-grid regularisation for Poisson inverse problems

Marta Lazzaretti (*Università degli studi di Genova/ Université Côte d’Azur*)

Fluorescence microscopy is a fundamental tool to investigate biological structures. Acquisitions are affected by blurring, arising from light diffraction, and noise, making the reconstruction of fine scale details challenging. In the recent years, several variational optimisation methods have been formulated in the continuous setting of measures spaces. Off-the-grid approaches are endowed with a solid mathematical theory and very good performance in applications for reconstructing point sources. Most of these approaches consider an ℓ^2 data term usually coupled with the TV norm (BLASSO problem) of the unknown measure. For better modeling the presence of a Poisson photon-counting noise process, we consider a non-quadratic data fidelity term, i.e. the Kullback-Leibler divergence and a non-negativity constraint. For the numerical computation of the solution, we adapt the Sliding Frank Wolfe algorithm to such scenario and perform a grid-search strategy for the regularisation parameter to compare the results obtained with the two data terms. We then propose an homotopy strategy for the automatic selection of the regularisation parameter, provided that an estimation of the noise level is known, which gives good quality reconstructions and significantly speeds up the algorithm. We present some comparative results in 1D and 2D and show a 3D real data reconstruction obtained with Homotopy-SFW with the KL.

Towards optimal sensor placement for inverse problems in spaces of measures

Daniel Walter (*Humboldt Universität zu Berlin*)

In this talk, we study the identification of a linear combination of point sources from a finite number of measurements contaminated by random noise. It relies on two main ingredients, first, a convex but non-smooth Tikhonov point estimator over the space of Radon measures and, second, a suitable mean-squared error based on its Hellinger-Kantorovich distance to the ground truth. Assuming standard non-degenerate source conditions as well as applying careful linearization arguments, a computable upper bound on the latter is derived. On the one hand, this allows to derive asymptotic convergence results for the mean-squared error of the estimator in the small variance case. On the other, it paves the way for applying optimal sensor placement approaches to sparse inverse problems.

Super-resolved Lasso

Clarice Poon (*Warwick University*)

Super-resolution of pointwise sources is of utmost importance in various areas of imaging sciences. Specific instances of this problem arise in single molecule fluorescence, spike sorting in neuroscience, astrophysical

imaging, radar imaging, and nuclear resonance imaging. In all these applications, the Lasso method (also known as Basis Pursuit or l_1 -regularization) is the de facto baseline method for recovering sparse vectors from low-resolution measurements. This approach requires discretization of the domain, which leads to quantization artifacts and consequently, an overestimation of the number of sources. While grid-less methods, such as Prony-type methods or non-convex optimization over the source position, can mitigate this, the Lasso remains a strong baseline due to its versatility and simplicity. In this work, we introduce a simple extension of the Lasso, termed ‘super-resolved Lasso’ (SR-Lasso). Inspired by the Continuous Basis Pursuit (C-BP) method, our approach introduces an extra parameter to account for the shift of the sources between grid locations. Our method is more comprehensive than C-BP, accommodating both arbitrary real-valued or complex-valued sources. Furthermore, it can be solved similarly to the Lasso as it boils down to solving a group-Lasso problem. A notable advantage of SR-Lasso is its theoretical properties, akin to grid-less methods. Given a separation condition on the sources and a restriction on the shift magnitude outside the grid, SR-Lasso precisely estimates the correct number of sources.

This is joint work with Gabriel Peyré.

TBA

Karl Hajjar (*Université Paris-Saclay*)

TBA

Flagfolds: multi-dimensional varifolds to handle discrete surfaces

Blanche Buet (*Université Paris-Saclay*)

We propose a natural framework for the study of surfaces and their different discretizations based on varifolds. Varifolds have been introduced by Almgren to carry out the study of minimal surfaces. Though mainly used in the context of rectifiable sets, they turn out to be well suited to the study of discrete type objects as well. While the structure of varifold is flexible enough to adapt to both regular and discrete objects, it allows to define variational notions of mean curvature and second fundamental form based on the divergence theorem. Thanks to a regularization of these weak formulations, we propose a notion of discrete curvature (actually a family of discrete curvatures associated with a regularization scale) relying only on the varifold structure. We performed numerical computations of mean curvature and Gaussian curvature on point clouds in $\mathbb{R}^{\#}$ to illustrate this approach. Though flexible, varifolds require the knowledge of the dimension of the shape to be considered. By interpreting the product of the Principal Component Analysis, that is the covariance matrix, as a sequence of nested subspaces naturally coming with weights according to the level of approximation they provide, we are able to embed all d -dimensional Grassmannians into a stratified space of covariance matrices. Building upon the proposed embedding of Grassmannians into the space of covariance matrices, we generalize the concept of varifolds to what we call flagfolds in order to model multi-dimensional shapes.

Joint work with: Gian Paolo Leonardi (Trento), Simon Masnou (Lyon) and Xavier Pennec (INRIA Sophia).

Exact recovery of the support of piecewise constant images via total variation regularization

Romain Petit (*Università degli studi di Genova*)

In this talk, I will consider the reconstruction of some unknown image from noisy linear measurements using total (gradient) variation regularization. Empirical evidence and theoretical results suggest that this method is particularly well suited to recover piecewise constant images. It is therefore natural to investigate how it performs when the unknown image has precisely this structure. I will present a noise robustness result stating that, in a low noise regime, the reconstruction is also piecewise constant, and one exactly recovers the number of shapes in the unknown image.

This is a joint work with Yohann De Castro and Vincent Duval.