Inverse Problems in Imaging
Foundations & Applications

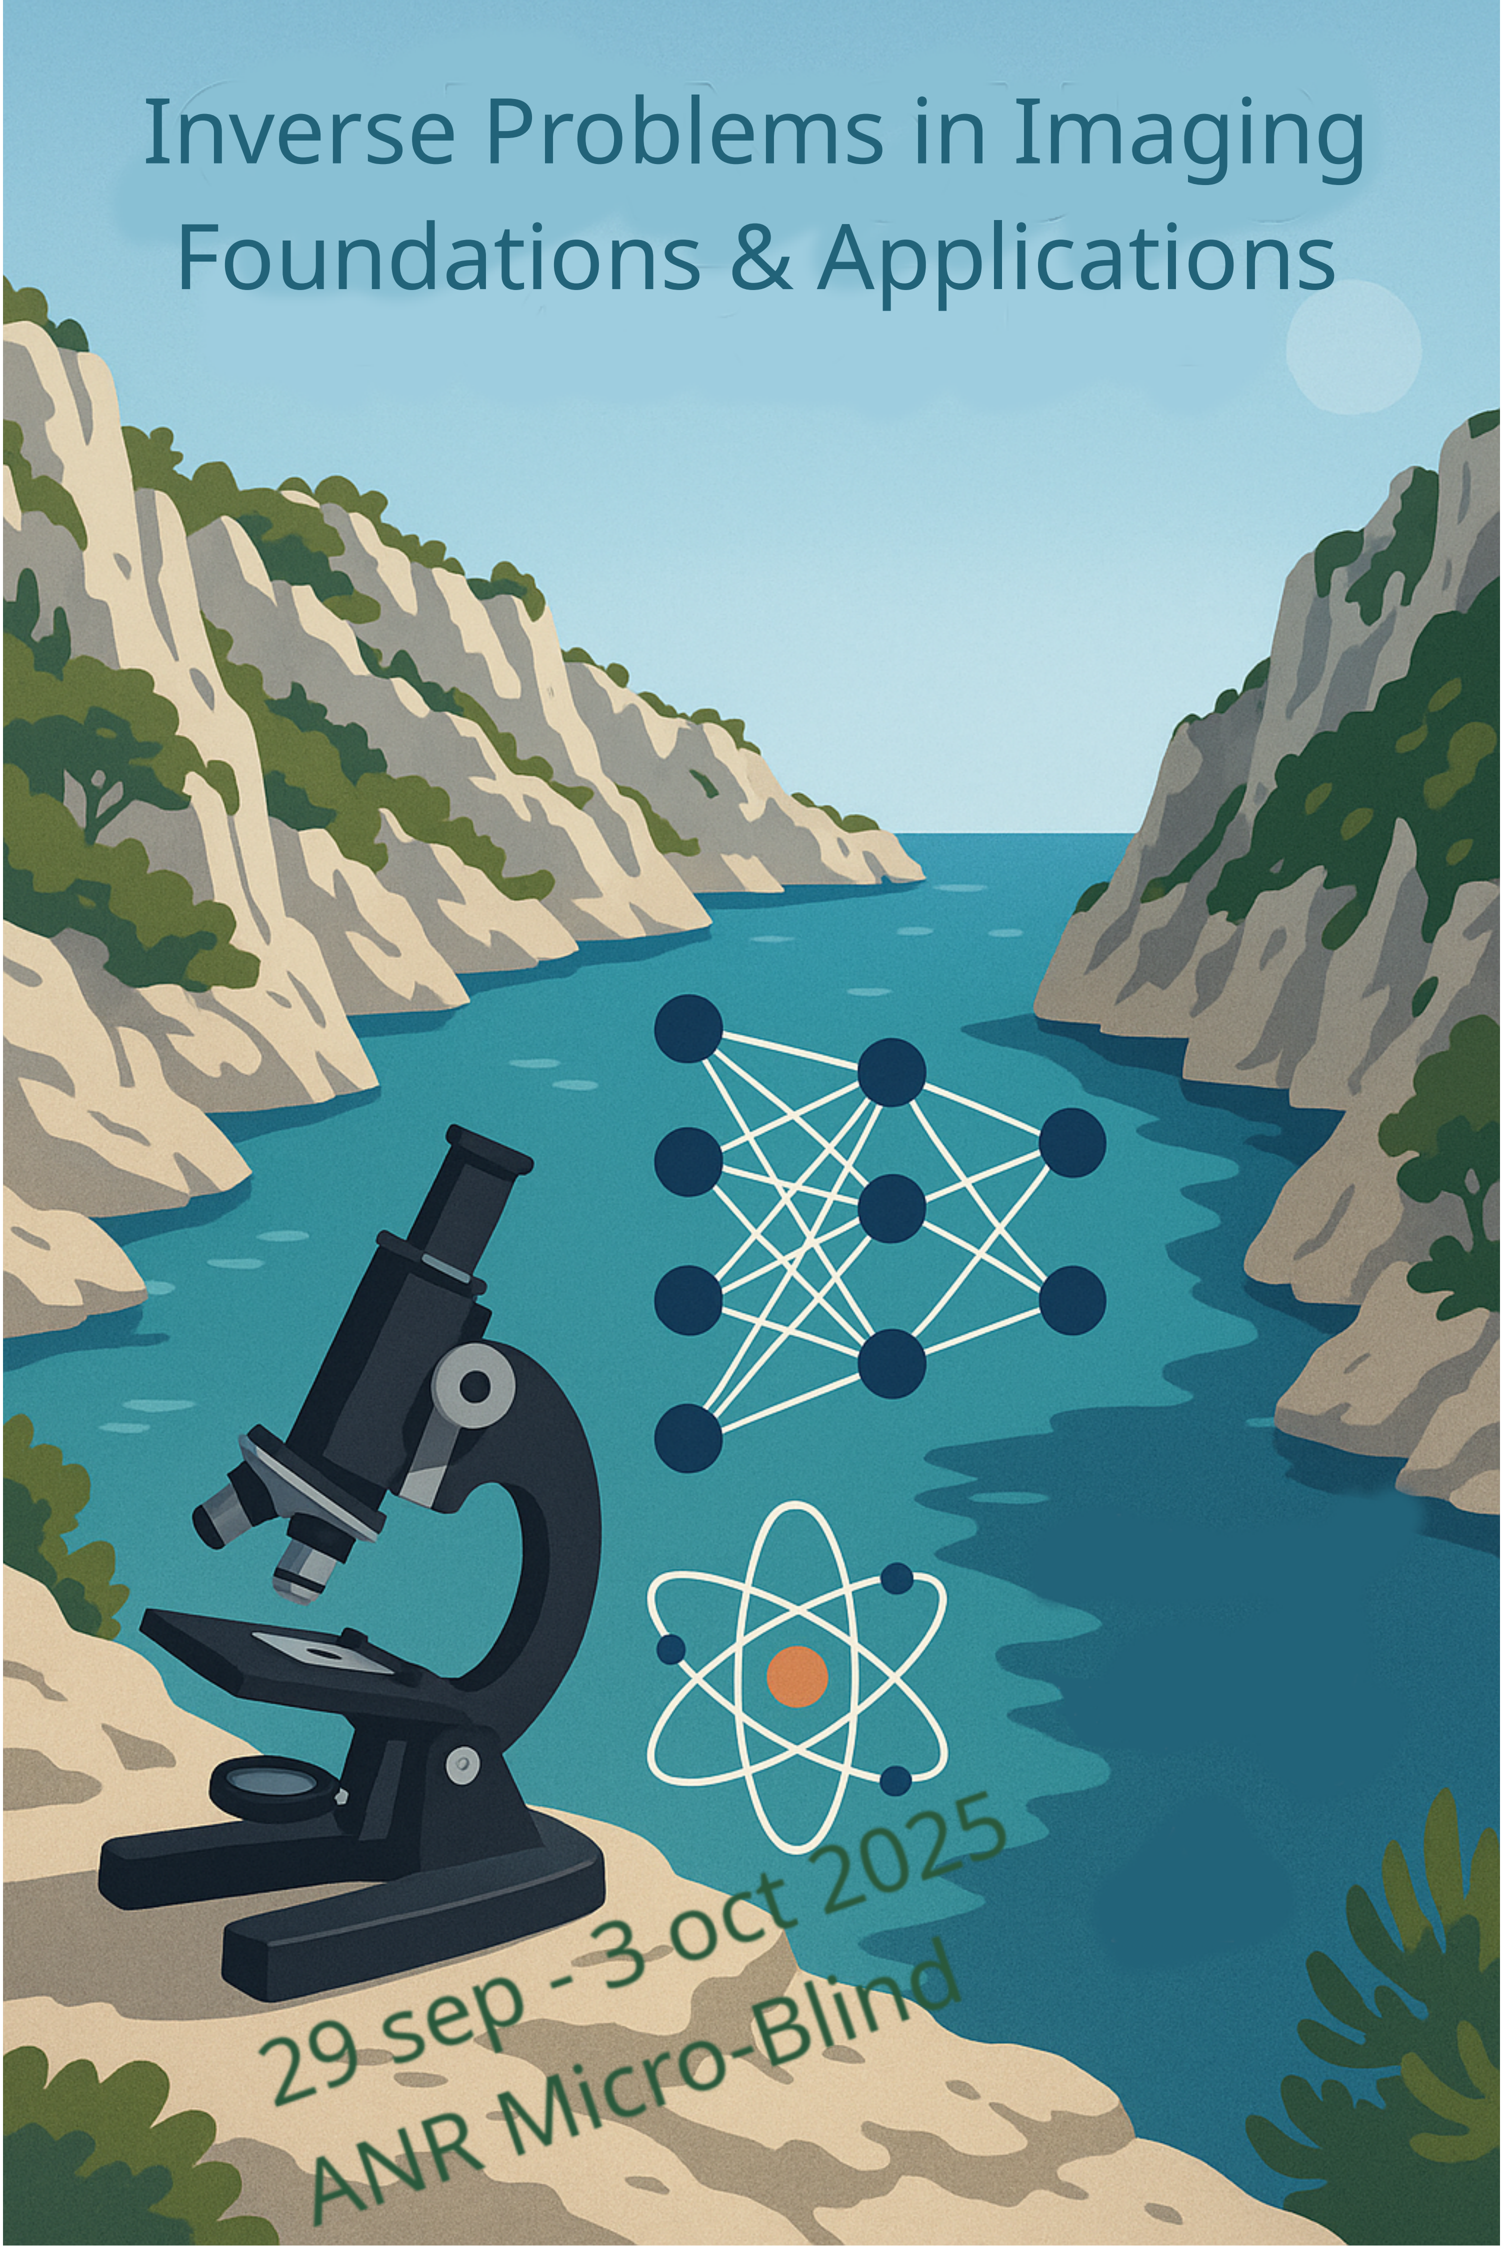

29 sep - 3 oct 2025

ANR Micro-Blind

The open-source LaTeX template, AMCOS_booklet, used to generate this booklet is available at https://github.com/maximelucas/AMCOS_booklet

Contents

About

Starting words

The conference *(Blind) inverse problems in imaging: from foundations to applications* is organized to celebrate the end of the Micro-blind project funded by the ANR and managed by the organizing committee.

The aim of the project is to develop new theories and algorithms to solve blind inverse problems and to apply them to the field of optical microscopy. The conference is an opportunity to gather specialists on the topic, present some of the results obtained during the project and prepare the forth-coming ones: the ERC JCJC TASKABILE *Task-adapted bilevel learning of flexible statistical models for imaging and vision* headed by Luca Calatroni and the ANR CLEAR-MICROSCOPY *Computational Learning for Efficient and Accurate Reconstruction in Microscopy* headed by Pierre Weiss.

The project members come from four laboratories:

- IMT (Institut de Mathématiques de Toulouse)
- IRIT (Institut de Recherche en Informatique de Toulouse)
- I3S (Laboratory of Computer Science, Signals and Systems of Sophia Antipolis)
- CBI (Centre de Biologie Intégrative, Toulouse)

The organizing committee[1],

Luca Calatroni
Emmanuel Soubies
Pierre Weiss

[1]Some last minute modifications of the program might occur due to unexpected events

Sessions planning

Monday 29/09	Tuesday 30/09	Wednesday 01/10	Thursday 02/10	Friday 03/10	
	Sampling	Super-resolution microscopy	Coupling biophysics & microscopy	PSF retrieval & engineering	1st morning session 9:00-10:20
Welcome	Coffee Break	Coffee Break	Coffee Break	Coffee Break	Coffee Break 10:20-10:50
Basics of Neural Network Solvers	Optimization for Learning	Computational Imaging	Network Training	Image Scanning Microscopy	2nd morning Session 10:50-12:30
Lunch Break	Lunch Break	Lunch Break	Lunch Break	Lunch Break	Lunch Break 12:30-14:00
Priors and Posterior	Mainfold & Inverse Problems	Free Time Calanques Pétanque	Diffusion again for those who missed it	Goodbye!	1 afternoon session 14:00-15:20
Coffee Break	Coffee Break		Coffee Break		Coffee Break 15:20-15:50
Unrolled methods	Operators & Tomography		Operator Learning		2nd afternoon session 15:50-16:50
					Free discussions
Discussion & Relax 17:30-19:30	Posters&Wine 17:30-19:30		Posters&Wine 17:30-19:30		Posters&Relax 17:30-19:30
Dinner	Dinner	Dinner	Dinner	Dinner	Dinner 19:30-21:00

	Monday 29/09	Tuesday 30/09	Wednesday 01/10	Thursday 02/10	Friday 03/10
1st morning session 9:00-10:20		Andres Almansa (40') Accelerating sampling Giacomo Meanti (20') Unsupervised diffusion Nicolas Papadakis (20') Posterior Sampling	Emmanuel Soubies (20') Intro super-resolution Anne Sentenac (30') RIM: theory Thomas Mangeat (30') RIM: practice	Hervé Turlier (40') Microscopy&Biophysics Fabian Erdel (20') Condensates modeling Oihan Joyot (20') Quantitative HALF-FRAP	Minh Hai Nguyen (20') Blind deconvolution Tobias Liaudat (20') Wavefront recovery Florian Sarron (20') PSF Identification
Coffee Break 10:20-10:50	Welcome	Coffee Break	Coffee Break	Coffee Break	Coffee Break
2nd morning Session 10:50-12:30	Pierre Weiss (20') IP Analytical Expressions Julie Delon (40') Intro Flow Matching Massias Mathurin (40') Expectation vs reality	Jalal Fadili (40') Inertial non convex Hippolyte Labarrière (20') Reparameterization Christian Daniele (20') Deep equilibrium Charles Dossal (20') Inertia	Hilton de Aguiar (40') Single pixel Pauline Trouvé (40') Privacy camera Amol Mahurkar (20') Energy Sensing	Rémi Gribonval (40') Dynamics of training Felix Krahmer (40') Convergence guarantees Samuel Vaiter (20') Properties of Autodiff	Lisa Cuneo & Luca Calatroni (60') ISM deconvolution Ending Words
Lunch Break 12:30-14:00	Lunch Break	Lunch Break	Lunch Break	Lunch Break	Lunch Break
1 afternoon session 14:00-15:20	Ulugbek Kamilov (40') Implicit priors Thomas Moreau (40') Conditional Probabilities Nathan Buskulic (20') Deep Inverse Priors	Paula Causin (40') Manifold & IPs Nathanael Munier (20') Adversarial Manifolds Samuel Hurault (20') Score & diffusion	Free Time Calanques Pétanque	Ségolène Martin Anne Gagneux (40') Diffusion models Mathieu Terris (20') Reconstruct Anything Model Hertrich Johannes (20') JKO sampling	Goodbye!
Coffee Break 15:20-15:50	Coffee Break	Coffee Break		Coffee Break	
2nd afternoon session 15:50-16:50	Audrey Repetti (40') Unfolded Proximal Caroline Chaux (20') Unrolling	Laurent Jacques (20') Implicit Spherical data Paul Escande (20') Photoacoustic tomo Valentin Debarnot (20') 3D tomography		Hadrien Montanelli (20') Inverse Scattering Romain Petit (20') Electrical Impedance Laure Blanc-Féraud (20') Curve reconstruction	
Free discussions	Free Time	Free Time		Free Time	
Posters&Relax 17:30-19:30		Posters&Wine 17:30-19:30		Posters&Wine 17:30-19:30	
Dinner 19:30-21:00	Dinner	Dinner	Dinner	Dinner Bouillabaisse	Dinner

List of Abstracts – Talks

Monday

The organizing committee, 10:20-10:40
Welcoming words
We will present how and why this conference was organized.

Pierre Weiss, 10:40-11:10
Analytical solutions for CNN inverse problem solvers
We provide analytical formulas for minimum mean square error estimators targeted at solving linear inverse problems, subject to constraints such as tranlastion equivariance and locality. This simple model turns out to predict surprisingly well the output of convolutional neural networks, at least on points close to the training set. This theory provides a rather clear path to studying facts such as how to inform the CNNs by the physics for the best performance.

Julie Delon, 11:10-11:50
From distributions to flows: an introduction to Flow Matching
Deep generative models aim to learn high-dimensional data distributions by transforming easily sampled reference measures into complex targets. Flow matching provides a principled way to construct such transformations by seeking vector fields that connect probability distributions through the continuity equation. In this talk, we will introduce the mathematical ideas behind flow matching and explore its links to score matching and diffusion models. We will also highlight the fruitful connections between flow matching and optimal transport.

Mathurin Massias, 11:50-12:30
On the Closed-Form of Flow Matching: Generalization Does Not Arise from Target Stochasticity
Modern deep generative models can now produce high-quality synthetic samples that are often indistinguishable from real training data. A growing body of research aims to understand why recent methods - such as diffusion and flow matching techniques - generalize so effectively. Among the proposed explanations are the inductive biases of deep learning architectures and the stochastic nature of the conditional flow matching loss. In this work, we rule out the latter – the noisy nature of the loss – as a primary contributor to generalization in flow matching. First, we empirically show that in high-

dimensional settings, the stochastic and closed-form versions of the flow matching loss yield nearly equivalent losses. Then, using state-of-the-art flow matching models on standard image datasets, we demonstrate that both variants achieve comparable statistical performance, with the surprising observation that using the closed-form can even improve performance.

Ulugbek Kamilov, 14:00-14:40
Computational Imaging: Restoration Deep Networks as Implicit Priors
Computational imaging problems are often formulated as ill-posed inverse problems, requiring the integration of prior knowledge to recover high-quality images from limited or corrupted measurements. This talk focuses on score-based methods, which approximate the gradient of the log-prior (known as the score function) using pre-trained deep restoration networks. I will present recent advances for learning score functions without access to clean training data. These include characterizing score functions via general restoration networks, learning from partial measurements, and model-based strategies that incorporate forward models into training. The talk will also cover the theoretical foundations of these approaches and their applications in biomedical image reconstruction.

Thomas Moreau, 14:40-15:20
Filling the gaps: a story of priors and conditional probabilities
Inverse problems are often ill-posed: incomplete measurements, noise, and ambiguity mean that we must rely on priors to reconstruct meaningful solutions. In the first part of this talk, I will present recent work that analyzes the limitations and challenges of these approaches when solving unsupervised inverse problems—how the choice of priors can bias reconstructions, how conditional models interact with measurement operators, and what this reveals about the fundamental difficulty of filling in missing information.

Building on this perspective, I will then introduce FIRE (Fixed point Restoration), a framework that addresses these challenges by defining implicit priors not just through denoisers but through general restoration models. The key idea is to characterize natural signals as fixed points of a degradation–restoration cycle, enabling a principled and flexible way to integrate pretrained networks into inverse problem solvers. This fixed-point view not only broadens the class of usable priors but also leads to robust, algorithms with strong empirical performance.

Nathan Buskulic, 15:20-15:40
Learning the optimal Tikhonov regularization for blind inverse problem, what should you expect ?
I will present in this talk how to extend the theory of optimal linear estimators, in the sense of expected mean-square error, from the non-blind setting to the blind setting where the interaction between the uncertainty of the operator and the signal interacts clearly. We will also see how the optimal linear estimator is in fact the solution of a generalized

Tikhonov regularized variational problem, and then we will discuss how to extend the classical Hölder source condition to the blind setting, and how it allows us to obtain new error bounds that are a generalization of the classical non-blind ones.

Audrey Repetti, 15:50-16:30
Analysis and synthesis approximated denoisers for forward-backward plug-and-play algorithms
In this presentation we will study the behaviour of the forward-backward (FB) algorithm when the proximity operator is replaced by a sub-iterative procedure to approximate a Gaussian denoiser, in a Plug-and-Play (PnP) fashion. Specifically, we consider both analysis and synthesis Gaussian denoisers within a dictionary framework, obtained by unrolling dual-FB iterations or FB iterations, respectively. We analyse the associated global minimization problems as well as asymptotic behaviour of the resulting FB-PnP iterations. For each case, analysis and synthesis, we show that the FB-PnP algorithms solve the same problem whether we use only one or an infinite number of sub-iteration to solve the denoising problem at each iteration. We will illustrate our theoretical results on numerical simulations, considering an image restoration problem in a deep dictionary framework. Joint work with Matthieu Kowalski, Benoit Malezieux and Thomas Moreau.

Caroline Chaux, 16:30-16:50
Learning Weighted Least Squares Through Unrolling for Poisson Image Deconvolution
We developed deep unroll networks to solve inverse problems in signal and image processing. The majority of recent works on algorithm unrolling has focused on learning the image prior (i.e. the regularisation term). However, the data fidelity term should also be carefully designed to account for the various degradations affecting the acquired data (e.g. noise, quantization, background signal). We explored the idea of learning a weighted least squares data fidelity term to better adapt to various types of noise in inverse problems. To achieve this, we proposed an end-to-end training approach combining a weight estimation module with an unrolled Forward-Backward Splitting algorithm where both the step size and regularization hyperparameters are learned. This results in a lightweight and highly interpretable architecture. We evaluated its performance on a classical Poisson image deconvolution task using a sparse wavelet-based prior, and numerical experiments demonstrate the advantages of learning a weighted least squares data term.

Abhijit Singh, Emmanuel Soubies, Caroline Chaux, *Learning Weighted Least Squares Data Term for Poisson Image Deconvolution*, ICASSP, Hyderabad, India, 6-11 Apr. 2025. `https://hal.science/hal-04887464`

Tuesday

Andres Almansa, 9:00-9:40
Accelerating Posterior Sampling with Generative Priors for Blind Inverse Problems
Posterior sampling is a key element when solving blind inverse problems via marginal likelihood maximization. Motivated by a blind deblurring application, we review the evolution of generative image priors, and the conditioning mechanisms (a.k.a. zero-shot PnP frameworks) that are required to turn such pretrained generic samplers into posterior samplers for a specific inverse problem. As generative models became more computationally efficient, the design of good conditioning mechanisms became also more sophisticated.

We finish the talk with a focus on Latent Consistency Models (LCMs), which distill latent-space text-to-image diffusion models (LDMs) into fast prior samplers. We leverage a new conditioning mechanism to propose LAtent consisTency INverse sOlver (LATINO), the first zero-shot PnP framework to solve inverse problems with priors encoded by LCMs. Our conditioning mechanism avoids automatic differentiation and reaches SOTA quality in as little as 8 neural function evaluations. We finally discuss the use of text conditioning in two contexts : *(i)* constraining the prior, when prior information is available, and *(ii)* blindly estimating the text prompt from the observation when prior information is not available.

Joint work with: Alessio Spagnoletti, Jean Prost, Nicolas Papadakis, Marcelo Pereyra, Charles Laroche, Eva Coupeté, Rémy Laumont, Valentin de Bortoli, Julie Delon, Alain Durmus.

Giacomo Meanti, 9:40-10:00
Unsupervised Imaging Inverse Problems with Diffusion Distribution Matching
I will talk about image restoration tasks addressed through the lens of inverse problems, having access to only unpaired datasets. The proposed method operates under minimal knowledge of the forward model, and no access to paired degraded and ground-truth images. This makes it particularly well-suited for real-world scenarios, where the forward model is often unknown or misspecified, and collecting paired data is costly or infeasible. The method leverages conditional flow matching to model the distribution of degraded observations, while simultaneously learning the forward model via a distribution-matching loss that arises naturally from the framework.

Nicolas Papadakis, 10:00-10:20
Posterior Sampling with the Proximal Stochastic Gradient Langevin Algorithm
We consider the problem of sampling distributions stemming from non-convex potentials with Unadjusted Langevin Algorithm (ULA). We prove the stability of the discrete-time

ULA to drift approximations under the assumption that the potential is strongly convex at infinity. In many context, e.g. imaging inverse problems, potentials are non-convex and non-smooth. Proximal Stochastic Gradient Langevin Algorithm (PSGLA) is a popular algorithm to handle such potentials. It combines the forward-backward optimization algorithm with a ULA step. Our main stability result combined with properties of the Moreau envelope allows us to derive a proof of convergence of the PSGLA for non-convex potentials. We empirically validate our methodology on synthetic data and in the context of imaging inverse problems. In particular, we observe that PSGLA exhibits faster convergence rates than Stochastic Gradient Langevin Algorithm for posterior sampling while preserving its restoration properties.

This is a joint work with Marien Renaud (IMB, Bordeaux), Valentin de Bortoli (Google DeepMind) and Artur Leclaire (Télécom Paris).

Fadili, Jalal, 10:50-11:30

Inertial Algorithms Meet NN-Based Methods for Inverse Problems

In this talk, I will focus on non-convex minimization problems via inertial second-order (in-time) dynamics and how they can prove valuable when solving inverse problems with neural network-based methods. I will first discuss several theoretical and practical issues for these algorithms, including convergence, convergence rates and trap avoidance properties. I will then turn to discussing how to bridge the worlds of optimization and that of inverse problems to provide convergence and recovery guarantees for a class of neural network-based methods (DeepInvese). I will also a precise characterization of the network architecture to benefit from these guarantees. This provides a first step towards the theoretical understanding of the interplay between the optimization dynamics and neural networks in the inverse problem setting.

Hippolyte Labarrière, 11:30-11:50

Reparameterization and Its Role in Optimization Dynamics

Recent advances in machine learning have been driven by increasingly overparameterized models, from deep neural networks to large language models. Combined with powerful computational infrastructure, these models now underpin many AI systems that have become part of everyday life. While their success is often attributed to massive training datasets and model scale, an equally important factor lies in how overparameterization shapes optimization and generalization. In this talk, we examine the impact of reparameterization on optimization procedures, highlighting its influence on implicit bias and trajectory convergence. By analyzing simplified models, we aim to build intuition for why reparameterization can accelerate learning and enhance generalization.

Christian Daniele, 11:50-12:10

Deep Equilibrium Models for Poisson Inverse Problems via Mirror Descent

Abstract: Inverse problems in imaging arise in a wide range of scientific and engineering applications, including medical imaging, astrophysics, and microscopy. These problems are inherently ill-posed, requiring advanced regularization techniques and optimization strategies to achieve stable and accurate reconstructions. In recent years, hybrid approaches that combine deep learning and variational methods have gained increasing attention. Well-established techniques include Algorithmic Unrolling, Plug-and-Play methods, and Deep Equilibrium Models. The latter are networks with fixed points, which are trained to match data samples from a training dataset. In this work, we focus on Deep Equilibrium Models to learn a data-driven regularization function for Poisson inverse problems, using the Kullback-Leibler divergence as the data fidelity term. To effectively handle this fidelity term, we employ Mirror Descent as the underlying optimization algorithm. We discuss theoretical guarantees of convergence, even in non-convex settings, incorporating a backtracking strategy, along with key aspects of training this class of models. To validate our approach, we evaluate its performance on a deblurring task with different kernels and varying levels of Poisson noise. Authors: Luca Calatroni, Silvia Villa, Samuel Vaiter, Christian Daniele.

Charles Dossal, 12:10-12:30
Inertia as a preconditioner
The primary reason inertial algorithms are used is their superior convergence speed relative to the number of iterations; typically, the convergence rate of gradient descent $O(1/n)$ can be improved to $O(1/n^2)$. However, it is less well known that these algorithms, when properly calibrated, can achieve a better convergence rate relative to the mu/L conditioning of the functional to be minimized.

Paola Causin, 14:00-14:40
Manifold Learning Approaches via Riemannian Geometry: Application to Inverse Problems in Biomedical Imaging
Biomedical image processing tasks can often be cast as inverse problems, where the measurement y arises from an ill-posed or ill-conditioned operator $G(x)$ perturbed by additive noise. The direct recovery of the desired field x from y is generally intractable, thereby necessitating regularization strategies that incorporate prior information. A widely adopted paradigm is to minimize a data fidelity functional $F(x)$ in conjunction with a regularization term $R(x)$, the latter being either prescribed *a-priori* or adaptively inferred by means of a neural network.

In this work, we propose a strategy where x is constrained to lie on a low-dimensional Riemannian manifold. To this end, we employ β-VAEs to derive a latent representation of the data. The pullback metric induced by the decoder is subsequently utilized to characterize the geometry of the manifold, with particular emphasis on the estimation of its intrinsic, a priori unknown, dimension.

In order to accommodate complex manifold topologies, potentially featuring discontinuities and nontrivial structures, we further incorporate a downstream Gaussian mixture model of local charts. Within this framework, each chart encodes a localized region of the manifold, whereas the mixture weights provide a principled representation of the global structure.

We assess the accuracy of intrinsic dimension estimation across a variety of benchmark manifolds and demonstrate the method applicability to CT reconstruction.

- Causin, P., Marta, A. *Estimating Dataset Dimension via Singular Metrics under the Manifold Hypothesis: Application to Inverse Problems.* arXiv 2507.07291, 2025.

- Alberti, G. S., Hertrich, J., Santacesaria, M., Sciutto, S. *Estimating Dataset Dimension via Singular Metrics under the Manifold Hypothesis: Application to Inverse Problems.* Journal of Machine Learning Research, 25(202), 1-35, 2024.

Nathanaël Munier, 14:40-15:00
Jackpot: Approximating Uncertainty Domains with Adversarial Manifolds
Given a forward mapping $\Phi : \mathbb{R}^N \to \mathbb{R}^M$, the region $\{x \in \mathbb{R}^N, \|\Phi(x) - y\| \leq \epsilon\}$, where $y \in \mathbb{R}^M$ is a given vector and $\epsilon \geq 0$ is a perturbation amplitude, represents the set of all possible inputs x that could have produced the measurement y within an acceptable error margin. This set reflects the inherent uncertainty or indeterminacy in recovering the true input x solely from the noisy observation y, which is a key challenge in inverse problems.

In this work, we develop a numerical algorithm called Jackpot (Jacobian Kernel Projection Optimization) which approximates this set with a low-dimensional adversarial manifold. The proposed algorithm leverages automatic differentation, allowing it to handle complex, high dimensional mappings such as those found when dealing with dynamical systems or neural networks. We demonstrate the effectiveness of our algorithm on various challenging large-scale, non-linear problems including parameter identification in dynamical systems and blind image deblurring.

Samuel Hurault, 15:20-15:40
From Denoising to Diffusion:
A Fine-Grained Error Analysis.
Sampling from an unknown distribution, accessible only through discrete samples, is a fundamental problem at the core of generative AI. The current state-of-the-art methods follow a two-step process: first, estimating the score function via denoising and then applying a diffusion-based sampling algorithm – such as Langevin or diffusion models. The resulting distribution's correctness can be impacted by four major factors: the generalization and optimization errors in deoising score matching, and the discretization and minimal noise amplitude in the diffusion. In this work, we provide a sharp analysis of the Wasserstein sampling error that arises from these four error sources. The four error

terms are made explicit in the Gaussian setting so as to exactly track how the anisotropy of the data distribution (encoded by its power spectrum) interacts with key parameters of the end-to-end sampling method. This result provides a foundation for further analysis of the tradeoffs involved in optimizing sampling accuracy.

Laurent Jacques, 15:50-16:10

Herglotz-NET: Implicit Neural Representation of Spherical Data with Harmonic Positional Encoding

Representing and processing data in spherical domains presents unique challenges, primarily due to the curvature of the domain, which complicates the application of classical Euclidean techniques. Implicit neural representations (INRs) have emerged as a promising alternative for high-fidelity data representation; however, to effectively handle spherical domains, these methods must be adapted to the inherent geometry of the sphere to maintain both accuracy and stability. In this context, we propose Herglotz-NET (HNET), a novel INR architecture that employs a harmonic positional encoding based on complex Herglotz mappings. This encoding yields a well-posed representation on the sphere with interpretable and robust spectral properties. Moreover, we present a unified expressivity analysis showing that any spherical-based INR satisfying a mild condition exhibits a predictable spectral expansion that scales with network depth. Our results establish HNET as a scalable and flexible framework for accurate modeling of spherical data.

Paul Escande, 16:10-16:30

On the numerics of photoacoustic tomography

Photoacoustic tomography (PAT) is a recent imaging modality that enables in vivo imaging of the distribution of nanometric agents with submillimetric resolution, a major challenge in biomedical imaging. However, fully harnessing the potential of this modality currently depends on image reconstruction algorithms that require significant computational resources (memory + time).

After introducing the underlying physics of these systems, we will turn to the simulation of forward models, with particular attention to computational efficiency.

Valentin Debarnot, 16:30:16:50

Supervised learning in tomography with partially unknown forward operator and no ground-truth

Tomography is a popular technique that is widely used in medicine, biology, material science, etc. Cryo-electron tomography (cryo-ET) is such a recent example. It enables 3D visualization of cellular environments. However, accurate reconstruction of high-resolution volumes is challenging due to the low signal-to-noise ratio and limited sample tilts, resulting in a missing wedge of Fourier information.

Recent advances in self-supervised deep learning have significantly improved reconstruction quality by post-processing initial reconstructions. However, these approaches are computationally expensive, memory-intensive, and require retraining a neural network for each new observation.

End-to-end supervised learning is an appealing alternative, but it is limited by the lack of reference data in cryo-electron microscopy, and training on synthetic data often leads to overfitting and poor generalization to real data.

To address these challenges, we introduce CryoLithe, a local, memory-efficient reconstruction network that directly estimates the volume from an aligned projection, overcoming the suboptimal back projection post-processing. We demonstrate that leveraging transform-domain locality makes CryoLithe robust to distribution shifts, enabling effective supervised training and giving excellent results on real data in minutes—much faster than existing methods—without retraining or fine-tuning.

Wednesday

Emmanuel Soubies, 9:00-9:20
Introduction to super-resolution fluorescent microscopy
For more than two decades, fluorescent super-resolution microscopy techniques have opened up new perspectives for biological research. They fall within the broader field of computational imaging, an area where optics and numerical processing combines to overcome the resolution limits of conventional systems. In this talk, after recalling the resolution limit of conventional optical systems, we will provide an overview of the main concepts—combining optics and numerical methods—that underlie the most widely used super-resolution techniques. We will then focus on the case of structured illumination microscopy and its variants, highlighting the main implementation challenges and the strategies developed to overcome them.

Anne Sentenac, 9:20-9:50
Super-resolved fluorescence microscopy using random illuminations (RIM)
In this talk, we will present the basis of super-resolved microscopy using random illuminations (RIM). This technique consists of recording multiple images of the sample under different speckled illuminations (obtained by passing the laser through a diffuser). The super-resolved image is reconstructed from the variance of the speckled images using a variance-matching algorithm. We will show the applicability of the concept to several imaging configurations (one- or two-photon fluorescence microscopy, non-linear microscopy).

Thomas Mangeat, 9:50-10:30
RIM and PRIM for subcellular dynamics on tissue or medium content screening.
We describe variance-based stochastic super-resolution imaging for live cell imaging with super-resolution comparable to the best 3D SIM.

The method consists of processing multiple images of the sample under different illumination conditions. The autocorrelation function of these images is controlled. We show that in the case of speckle illumination,

where the speckle correlation length coincides with the width of the observation point spread function, there is clearly a twofold increase in resolution [2]. Using a variance-matching algorithm called AlgoRIM [2-3-4-5-6],

the super-resolved reconstruction is obtained numerically from the variance of the speckle images and the autocorrelation function of the speckles. This method allows the combination of strong optical sectioning and super-resolution

for biological samples with out-of-focus fluorescence. In addition, the method is not affected by optical aberrations on the excitation side, is linear with respect to brightness, and is compatible with multi-colour live cell imaging

over long time periods [2-3-6-9]. We will show a practical implementation from the TIRF configuration [2] to projected EDF-RIM imaging of large volumes in a single setup [6].

Finally, we will show that the scope of variance-based fluorescence imaging can be significantly extended by extending the concept to illumination other than speckle [7]. Live imaging from cell culture to tissue will be presented.

The methods improve 3D imaging and the reproducibility of media content analysis in complex conditions such as expansion microscopy or organoids [8], or the study of focal adhesion dynamics of human osteoclasts on bone [9].

Finally, the extension of long-term imaging allowed us to identify for the first time a novel chromatin domain D that regulates the response to DNA double-strand breaks [10].

1. Labouesse, S., Idier, J., Allain, M., Giroussens, G., Mangeat, T., & Sentenac, A. (2024). *Super-resolution capacity of variance-based fluorescence microscopies, from Random Illumination Microscopy (RIM) to Super-resolved Optical Fluctuation Imaging (SOFI).*Physical Review A.

2. Mangeat, T., Labouesse, S., Allain, M., Negash, A., Martin, E., Guénolé, A., ... & Sentenac, A. (2021). *Super-resolved live-cell imaging using Random Illumination Microscopy.* Cell Reports Methods, 1(1), 100009.

3. Affannoukoué, K., Labouesse, S., Maire, G., Gallais, L., Savatier, J., Allain, M., ... & Sentenac, A. (2023). *Super-resolved total internal reflection fluorescence microscopy using random illuminations.*Optica, 10(8), 1009-1017.

4. Giroussens, Guillaume, Labouesse, Simon, Allain, Marc, et al. *Fast super-resolved reconstructions in fluorescence random illumination microscopy (RIM).* IEEE Transactions on Computational Imaging, 2024.

5. `https://github.com/teamRIM/tutoRIM`

6. Mazzella, L., Mangeat, T., Giroussens, G., Rogez, B., Li, H., Creff, J., ... & LeGoff, L. (2024). *Extended-depth of field random illumination microscopy, EDF-RIM, provides super-resolved projective imaging. Light: Science & Applications*, 13(1), 285.

7. Pierre Barbault, Jérôme Idier, Simon Labouesse, Marc Allain, Thomas Mangeat, et al.. *Pseudo-Random Illumination Microscopy.* 2025. ⟨hal-04977967v1⟩

8. Schweizer, N., Haren, L., Dutto, I., Viais, R., Lacasa, C., Merdes, A., & Lüders, J. (2021). *Sub-centrosomal mapping identifies augmin-TuRC as part of a centriole-stabilizing scaffold.* Nature communications, 12(1), 1-16.

9. Portes, M., Mangeat, T., Escalier, N., Raynaud-Messina, B., Thibault, C., Maridonneau-Parini, I., ... & Poincloux, R. (2021). *Nanoscale architecture and coordination of actin cores within the sealing zone of human osteoclasts.* Elife

10. Arnould, Coline, Rocher, Vincent, Saur, Florian, et al. *Chromatin compartmentalization regulates the response to DNA damage.* Nature, 2023, p. 1-10.

Hilton B. de Aguiar, 10:50-11:30

Single-pixel methods for linear and nonlinear microscopy

Imaging complex media with high-resolution (sub-um) is a challenge with widespread applications in various industrial and fundamental research settings. However, complex media imaging presents various challenges such as high physical (e.g. refractive index) and chemical heterogeneity. In this presentation, I will introduce single-pixel methods that aim to (i) reach high-resolution deep imaging and (ii) high-speed chemically-selective imaging using the Raman effect. I will motivate the use of single-pixel detectors, compared to the classical multi-pixel approach, present their mathematical background and our proposed solutions.

Pauline Trouvé-Peloux, 11:30-12:10

End-to-end design of imaging systems and neural networks - application to the design of a privacy preserving camera

End-to-end imaging system design consists of the joint optimization of an optical system and a processing to maximize the quality of the outcome at the end of the processing chain. When the processing involves a neural network, a derivable optical simulation enables the simultaneous optimization of neural network parameters and those of the optics. We will present the principle behind this approach, the simulation tools developed at ONERA to implement it, and an application in the context of the end-to-end design of a privacy-preserving camera. Indeed, today, we are surrounded by cameras performing tasks like image content analysis, notably using neural networks (semantic segmentation, pose estimation, action detection). However, their use in urban areas is limited due to the risks they pose regarding privacy preservation. One method to offer better privacy guarantees involves using sufficiently degraded images to prevent individual recognition. Designing such a camera, whose image quality is intentionally degraded while maintaining effective analysis performance, is a typical end-to-end design challenge. We will outline the approach developed at ONERA to optimize this type of camera, relying on an adversarial learning between two neural networks: one dedicated to a useful task, that is expected to perform well, and one attacking privacy, that we are seeking to prevent, all this wile optimizing parameters of a phase mask. Experimental results obtained with the co-designed optics will be presented.

Amol Mahurkar, 12:10-12:30
Generalized Energy Sensing for Imaging
In this talk, I present a novel mathematical sensing framework for optical imaging modalities, among others, that leverage fast, photon-sensitive, non-spatially-resolved detectors such as photodiodes (PDs) and bucket detectors (PMTs), in conjunction with wavefront modulators such as digital micromirror devices (DMDs) and spatial light modulators (SLMs). We introduce a Generalized Energy Sensing for Imaging (GESI) framework as a generalization of the single-pixel camera model, enabling imaging through complex media. The GESI formulation generalizes energy sensing as optical modes become mixed when passing through complex media—specifically, by incoherently integrating spatially, temporally, or spectrally onto such detectors (PDs, PMTs). Next, we address the non-convexity challenges that GESI entails and propose a spectral method that recovers the complex-valued object exactly (up to a global phase), in expectation, providing a foundation for gradient-based methods.

Thursday

Hervé, Turlier, 9:00-9:40
Bridging Fluorescence Microscopy and Biophysical Tissue Models
Fluorescence microscopy is a key tool for studying biological systems, yet extracting physical insights from 3D images remains challenging. Meanwhile, tissue models are becoming increasingly sophisticated, but direct integration with imaging data is still limited. In this talk, I will present our recent efforts to bridge this gap. I will introduce our

computational foam-like tissue models, which incorporate viscous dissipation, cell division, and mechanochemical feedback. Then, I will present a segmentation and 3D tension inference method that generates detailed mechanical atlases of embryos and tissues from microscopy images. Finally, I will showcase a fully differentiable optimization pipeline that links mechanical models to microscopy by generating realistic synthetic images from simulations, paving the way for solving inverse mechanical problems.

Fabian Erdel, 9:40-10:00

Biomolecular condensates: Structure, dynamics and mechanics

Cells contain macromolecules, many of which are intrinsically disordered. These macromolecules tend to interact with each other to form so-called biomolecular condensates, which organize cellular activities in time and space. Condensates can form through different molecular mechanisms, giving rise to assemblies with different biophysical properties. In particular, the internal structure, the diffusive motion of molecules and the mechanical properties of condensates vary depending on the underlying assembly mechanism. I will illustrate the properties of cellular condensates along with the main mathematical and physical underpinnings, and I will discuss the current challenges and the biological consequences of condensate formation.

Oihan Joyot, 10:00-10:20

Learning the Dynamic Law of Cellular Condensates

Cellular condensates regulate biochemical activities in time and space, but their dynamics remain poorly understood, even though dysfunctions can lead to severe diseases such as cancer and neurodegenerative disorders. In this work, we propose a technique for inferring key parameters such as condensate permeability and viscosity. This approach is based on FRAP microscopy, which enables the visualization of the return to equilibrium of fluorescent molecules after photobleaching. We present a mathematical model in which fluorescent molecules are treated as particles whose positions and velocities follow stochastic motion laws, and we derive the corresponding PDEs to enable faster simulations using the finite difference method. The parameters of this model are then inferred through a regression problem applied to molecular dynamics videos. Preliminary encouraging results are presented, providing access to previously inaccessible quantities.

Rémi Gribonval, 10:50-11:30

Training dynamics of ReLU Networks: a Path-lifting Perspective

Can we hope to decipher the role of the well-known rescaling symmetries of ReLU networks parameterizations in their training dynamics ? The talk will explore recent advances in this direction that exploit the path-lifting, a rescaling-invariant polynomial representation of the parameters of general ReLU networks. Despite its combinatorial dimension, the path-lifting turns out to be not only a convenient mathematical analysis tool: it also gives rise to a computational toolbox to reveal useful properties of the function corresponding to a ReLU network, from Lipschitz regularity to convexity.

Felix Krahmer, 11:30-12:10

Solving Inverse Problems with Deep Linear Neural Networks: Global Convergence Guarantees for Gradient Descent with Weight Decay

Machine learning methods are commonly used to solve inverse problems, wherein an unknown signal must be estimated from few measurements generated via a known acquisition procedure. In particular, neural networks perform well empirically but have limited theoretical guarantees. In this work, we study an underdetermined linear inverse problem that admits several possible solution mappings. A standard remedy (e.g., in compressed sensing) establishing uniqueness of the solution mapping is to assume knowledge of latent low-dimensional structure in the source signal. We ask the following question: do deep neural networks adapt to this low-dimensional structure when trained by gradient descent with weight decay regularization? We prove that mildly overparameterized deep linear networks trained in this manner converge to an approximate solution that accurately solves the inverse problem while implicitly encoding latent subspace structure. To our knowledge, this is the first result to rigorously show that deep linear networks trained with weight decay automatically adapt to latent subspace structure in the data under practical stepsize and weight initialization schemes. Our work highlights that regularization and overparameterization improve generalization, while overparameterization also accelerates convergence during training. This is joint work wih Hannah Laus (Technical University of Munich) as well as Suzanna Parkinson (University of Chicago), Vasileios Charisopoulos (University of Washington), and Rebecca Willett (University of Chicago).

Samuel Vaiter, 12:10-12:30

Remarks about bilevel optimization

Despite the successes of bilevel optimization in hyperparameter optimization and metalearning, it commonly admitted that BO is complex from a theoretical and practical perspectives. In this talk, I will discuss several theoretical properties that highlight why BO is significantly harder than classical optimization.

Ségolène Martin & Anne Gagneux, 14:00-14:40

From flow matching to denoising and back

We start our talk introducting Plug-and-Play (PnP) Flow Matching, an algorithm for solving imaging inverse problems. PnP methods leverage powerful pre-trained denoisers, while Flow Matching (FM) has pushed the frontier of generative sampling. Our method combines both by defining a time-dependent denoiser from a pre-trained FM model, alternating data-fidelity updates, reprojections onto the FM path, and denoising. It is efficient and memory-friendly, and achieves superior results on diverse inverse problems compared to existing PnP and FM-based methods.

Building on this, we ask: is a "generative" denoiser nothing more than a very good denoiser? And are good generative models essentially nothing more than good denoisers across all noise levels? To address these questions, we establish formal links between denoisers and generative models, and introduce a denoising toolkit, i.e., controlled procedures to test the impact of several factors on the performance of flow matching models. We show that different denoising losses and parameterizations, though theoretically equivalent if perfectly trained, lead to very different empirical performance. Moreover, we demonstrate that different phases of the generative process (early, intermediate, and late) play distinct roles in determining sample quality. Our analysis highlights in particular the importance of the intermediate stage.

Matthieu Terris, 14:40-15:00
Reconstruct Anything Model: generalizing restoration models beyond a single task
Most existing learning-based methods for solving imaging inverse problems can be roughly divided into two classes: iterative algorithms, such as plug-and-play and diffusion methods, that leverage pretrained denoisers, and unrolled architectures that are trained end-to-end for specific imaging problems. Iterative methods in the first class are computationally costly and often provide suboptimal reconstruction performance, whereas unrolled architectures are generally specific to a single inverse problem and require expensive training. In this presentation, we propose a novel non-iterative, lightweight architecture that incorporates knowledge about the forward operator (acquisition physics and noise parameters) without relying on unrolling. Our model is trained to solve a wide range of inverse problems beyond denoising, including deblurring, magnetic resonance imaging, computed tomography, inpainting, and super-resolution. The proposed model can be easily adapted to unseen inverse problems or datasets with a few fine-tuning steps (up to a few images) in a self-supervised way, without groundtruth references. Throughout a series of experiments, we demonstrate state-of-the-art performance from medical imaging to low-photon imaging and microscopy.

Johannes Hertrich, 15:00-15:20
Importance Corrected Neural JKO Sampling
We sample from an unnormalized probability density function by iteratively training normalizing flows with a regularized loss function. We relate this sampling scheme with the JKO scheme for simulating Wasserstein gradient flows. In order to overcome local minima and slow convergence of the Wasserstein gradient flow, we combine it with importance-based rejection-resampling steps. We prove that the KL divergence to the target distribution decreases in each step and demonstrate the efficiency of the scheme by numerical examples. Finally, we extend our framework to importance-based fine-tuning of score-based diffusion models for downstream tasks like posterior sampling in inverse problems.

Montanelli Hadrien, 15:50-16:10
Neural networks for inverse scattering
We present a neural network framework for selecting regularization parameters in the linear sampling method for inverse acoustic scattering problems. Our approach combines

two complementary networks: a standard neural network that estimates the noise level, and a neural operator that chooses the corresponding regularization parameters. Both networks train quickly thanks to exact solutions on disks, yet they generalize well to more complex scatterer shapes and material configurations. Two-dimensional experiments confirm the accuracy of the method, and interpretability analyses offer insight into the networks' behavior.

Romain Petit, 16:10-16:30
On the non-convexity issue in electrical impedance tomography
A classical approach in electrical impedance tomography is to estimate the unknown conductivity by solving a nonlinear least-squares problem. It leads to a nonconvex optimization problem which is generally believed to be riddled with bad local minimums. We revisit this issue in the case of piecewise constant radial conductivities and prove that, contrary to previous claims, there are no local minimums in the case of two scalar unknowns with no measurement noise. We also provide a partial proof of this result in the general setting which holds under a numerically verifiable assumption. Finally, we investigate whether a recently proposed approach based on convexification yields better reconstructions. For the first time, we propose a way to implement it in practice and show that it is consistently outperformed by Newton-type least squares solvers, which are also faster and require less measurements.

Laure Blanc-Féraud, 16:20-16:50
Off-the-Grid Curve Reconstruction in blurred Images by optimization of functional.
Recently an optimization functional has been introduced for the reconstruction of off-grid curves in inverse problems in image. This functional is defined on the space of finite divergence Radon vector measures. Its numerical resolution presents major challenges. We will present the solutions we have explored to address these issues, showing theoretical results, sticking points, and numerical results. In particular theoretical results are only available for a linear data term, which we have shown is not well suited to curve reconstruction in a blurred image. Finally, we will propose a new promising way of writing the functional for the reconstruction of off-the-grid curves.

Friday

Minh Hai Nguyen, 9:00-9:20
How Diffusion Prior Landscapes Shape the Posterior in Blind Deconvolution
The Maximum A Posteriori (MAP) estimation is a widely used framework in blind deconvolution to recover sharp images from blurred observations. The estimated image and blur filter are defined as the maximizer of the posterior distribution. However,

when paired with sparsity-promoting image priors, MAP estimation has been shown to favors blurry solutions, limiting its effectiveness. In this paper, we revisit this result using diffusion-based priors, a class of models that capture realistic image distributions. Through an empirical examination of the prior's likelihood landscape, we uncover two key properties: first, blurry images tend to have higher likelihoods; second, the landscape contains numerous local minimizers that correspond to natural images. Building on these insights, we provide a theoretical analysis of the blind deblurring posterior. This reveals that the MAP estimator tends to produce sharp filters (close to the Dirac delta function) and blurry solutions. However local minimizers of the posterior, which can be obtained with gradient descent, correspond to realistic, natural images, effectively solving the blind deconvolution problem. Our findings suggest that overcoming MAP's limitations requires good local initialization to local minima in the posterior landscape. We validate our analysis with numerical experiments, demonstrating the practical implications of our insights for designing improved priors and optimization techniques.

Tobias Liaudat, 9:20-9:40

Point spread function wavefront recovery: phase retrieval with automatic differentiation
Upcoming space telescopes with wide-field optical instruments, like the Euclid space mission, have a spatially varying Point Spread Function (PSF). Modelling this PSF field is a challenging ill-posed inverse problem, as an accurate model must be built from a limited number of noisy, undersampled, and spectrally-integrated star observations. The PSF is governed by the system's underlying wavefront error (WFE), but recovering this phase information from standard in-focus images is a notoriously difficult task. Standard data-driven models build the PSF directly in the pixel space, which proved insufficient for the stringent performance requirements. Our previous work on data-driven models can generate accurate PSFs in the pixel space, but it fails to recover a physically meaningful WFE.

In this talk, I will present a novel phase retrieval method that successfully estimates the WFE by exploiting the spatial diversity of aberrations across a wide field of view. Our approach leverages a differentiable optical model that links all observations to a common WFE field. We introduce a new optimisation strategy that circumvents the direct, unstable optimisation of the physical model's parameters. Instead, we iteratively train a flexible, over-parameterised data-driven component and then project the learned information onto a physical basis of Zernike polynomials. This projection-reset loop, enabled by automatic differentiation, allows the model to escape poor local minima and progressively build an accurate estimate of the ground-truth WFE. We will demonstrate through simulations that this method recovers the underlying WFE field from only degraded, in-focus observations.

Florian Sarron, 9:40-10:00

Learning to identify PSFs in fluorescence microscopy

Accurate identification of the point spread function (PSF) is essential for improving the quality of reconstructions in fluorescence microscopy. We present a statistical learning–based approach to automatically estimate the PSF directly from acquired images, without relying on fluorescent beads or dedicated calibration procedures. I will discuss the performance of this method on both simulated and real datasets with known induced PSFs, highlighting its robustness to varying blur complexity and image noise levels.

Lisa Cuneo & Luca Calatroni, 10:50-11:30

Reconstruction Approaches in Image Scanning Microscopy: Regularization and Optimization

Image Scanning Microscopy (ISM) is an advanced fluorescence microscopy technique that improves spatial resolution by replacing a single "bucket" detector with an array of detectors. This change provides access to additional spatial information, producing a multi-dimensional dataset. When ISM is implemented with Single-Photon Avalanche Diode (SPAD) array detectors, further advantages emerge: they enable true single-photon time tagging (sub-nanosecond resolution) with no read-out noise. Given such ISM acquisitions, the problem of reconstructing an high-resolution specimen image naturally falls within the framework of inverse problems with Poisson data. Reconstruction is commonly tackled using the iterative Richardson–Lucy (RL) algorithm, which can be extended to a multi-image setting that reflects the acquisition process. However, RL shows amplification of noise as the iteration progresses. To avoid it, one typically incorporates prior information, such as Tikhonov regularization, ℓ^1 or total variation penalties, or even plug-and-play procedures embedded in suitable optimization schemes. While effective, these approaches rely on structural assumptions about the specimen. In contrast, we propose an alternative strategy which leverages the time-tagging detection to easily obtain multiple realizations of the same Poisson process: a variance-based regularization that encodes information from the acquisition process itself by penalizing discrepancies between independent realizations. Unlike structural priors, this "agnostic" prior is grounded in noise statistics. It effectively mitigates noise amplification and enhances reconstruction stability.

List of Participants

First name	Last name	Institution
Hilton Barbosa	Aguiar	CNRS MAP5 - Univ. Paris Cité
Andres	Almansa	CNRS Kastler Brossel/ENS Paris
Laure	Blanc-Féraud	CNRS, Univ. Côte d'Azur, I3S
Nathan	Buskulic	Univ. degli Studi di Genova
Luca	Calatroni	Univ. di Genova
Paola	Causin	Univ. of Milano
Caroline	Chaux	CNRS, Univ. Marseille, I2M
Jonathan-Eduardo	Chinos-Rodriguez	Univ. Toulouse, IRIT
Lisa	Cuneo	Instituto Italiano di Tecnologia
Christian	Daniele	Univ. of Genoa
Valentin	Debarnot	CREATIS, INSA Lyon
Fabrice	Delbary	CNRS Collège de France
Julie	Delon	Université Paris Cité
Charles	Dossal	Univ. Toulouse, IMT
Fabian	Erdel	CNRS, Univ. Toulouse, CBI
Paul	Escande	Univ. Toulouse, IMT
Mhamed	Essafri	Univ. Toulouse
Jalal	Fadili	ENSICAEN, CNRS
Anne	Gagneux	ENS de Lyon
Rémi	Gribonval	ENS de Lyon
Johannes	Hertrich	Univ. Paris Dauphine
Samuel	Hurault	ENS - PSL
Laurent	Jacques	UCLouvain
Oihan	Joyot	IRIT (toulouse)
Ulugbek	Kamilov	Washington Univ.
Felix	Krahmer	Technical Univ. of Munich
Hippolyte	Labarrière	Univ. Genoa
Tobias	Liaudat	CEA-Saclay
Amol	Mahurkar	Kastler Brossel, ENS, Paris
Thomas	Mangeat	Univ. Toulouse, CBI
Ségolène	Martin	TU Berlin
Mathurin	Massias	INRIA Lyon
Giacomo	Meanti	INRIA
Hadrien	Montanelli	INRIA Paris-Saclay
Thomas	Moreau	INRIA
Nathanaël	Munier	Univ. Toulouse, IRIT&CBI
Minh Hai	Nguyen	Univ. Toulouse, IRIT&CBI
Nicolas	Papadakis	CNRS Univ. Bordeaux, IMB
Romain	Petit	CNRS and ENS
Audrey	Repetti	Heriot-Watt University
Aude	Rondepierre	Univ. Toulouse, IMT
Florian	Sarron	Univ. Toulouse, IMT
Anne	Sentenac	CNRS, Institut Fresnel
Emmanuel	Soubies	CNRS, Univ. Toulouse, IRIT
Matthieu	Terris	INRIA
Pauline	Trouvé	ONERA
Hervé	Turlier	CNRS, Collège de France
Samuel	Vaiter	CNRS Univ. Côte d'Azur
Pierre	Weiss	CNRS IRIT& CBI

Useful Information

Talks will be held at the lecture room A1

Coffee breaks and lunches will be offered next to the room.

The **poster session** will be held on Tuesday and Thursday evening together with a reception.

The **conference dinner** will be held on Thursday at the CIRM restaurant with the world-famous CIRM Bouillabaisse.

How to get to CIRM

Detailed instructions are given there: `https://www.cirm-math.com/getting-to-cirm.html`.

From gare de Marseille Saint-Charles, count about 45 minutes.

- take the Metro 2 (Sainte-Marguerite Dromel) from Saint-Charles to Rond-Point du Prado (about 8 minutes)
- take bus B1 or 21J to the terminus, Luminy campus (about 20 minutes)
- Walk up (about 12 minutes)
- Check Google Maps if any problem
- If you are really, really lost and surrounded by wild boars, call +33 6 66 48 11 85 (P. Weiss's phone)