

Mathematics and
Image
Analysis

MIA **2025** Conference

Paris, 13 - 15 January, 2025
Conference Booklet

WELCOME

Welcome to the twelfth edition of the Mathematics and Image Analysis (MIA) conference series! The conference builds on a series of successful MIA conferences initiated in 2000, which have been biennially hosted at IHP in Paris. Since 2014, German researchers have joined in organizing the conference and it was decided to alternate its venue between Paris and Berlin.

The scientific program includes invited lectures at the interface between research in applied mathematics (PDEs, statistics, inverse problems, optimization, geometrical modeling, optimal transport, etc.) and new developments in various areas of imaging science, machine learning and high-dimensional data processing, related to topics including restoration, compressed sensing, natural image modeling, and neuro-imaging.

Scientific Committee

Martin Burger
Rémi Gribonval
Michael Hintermüller
Nelly Pustelnik
Carola-Bibiane Schönlieb
Silvia Villa

Organizing Committee

Leon Bungert
Julie Delon
Nicolas Papadakis
Gabriele Steidl
Pauline Tan
Samuel Vaiter

Conference Website

rt-maiages.math.cnrs.fr/mia25/mia25

CONTENTS

Schedule	3
Abstracts	4
Ph.D. Prize	12
Posters	13
Venue	20
Facilities	21

SPONSOR



Note: A printable version of this booklet is available on the website of the conference.

SCHEDULE-AT-A-GLANCE

13 MON	<i>Welcome message</i>	10:00 a.m.
	Silvia Bonnetini Linesearch-enhanced forward-backward methods for inexact nonconvex...	11:00 a.m.
	Johannes Hertrich Importance corrected neural JKO sampling	11:30 a.m.
	<i>Lunch break</i>	12:00 p.m.
	<i>Ph.D. prize</i>	2:00 p.m.
	Mame Diarra Fall Regularization by denoising: Bayesian model and Langevin-within-split...	2:30 p.m.
	Marcelo Pereyra Uncertainty quantification in statistical imaging sciences: 40 years of...	3:00 p.m.
	<i>Coffee break</i>	3:30 p.m.
	Benedikt Wirth The shape space of Sobolev diffeomorphisms and its discretization	4:00 p.m.
	Elena di Bernardino Curvature measures for random excursion sets: theoretical and...	4:30 p.m.
14 TUE	Audrey Repetti Analysis and synthesis approximated denoisers for forward-backward plug...	9:00 a.m.
	Barbara Pascal Bilevel optimization for automated data-driven inverse problem resolution	9:30 a.m.
	Elisa Ricietti Exploiting multiple resolutions to accelerate inverse problems in imaging	10:00 a.m.
	<i>Coffee break</i>	10:30 a.m.
	Kimia Nadjahi Scalable unbalanced optimal transport by slicing	11:00 a.m.
	Matthew Thorpe How many labels do you need in semi-supervised learning?	11:30 a.m.
	<i>Lunch break</i>	12:00 p.m.
	Antonin Chambolle Some properties of the solutions of Total-Variation regularized inverse...	2:00 p.m.
	Clarice Poon Hadamard Langevin dynamics for sampling sparse priors	2:30 p.m.
	Matthias J. Ehrhardt Inexact algorithms for bilevel learning	3:00 p.m.
<i>Coffee break</i>	3:30 p.m.	
<i>Poster session</i>	4:00 p.m.	
15 WES	Jonathan Dong Random Phase Retrieval: theory and implementation	9:00 a.m.
	Anne Wald Nano-CT imaging as an inverse problem with inexact forward operator	9:30 a.m.
	Luca Calatroni Deep equilibrium mirror descent for learning image regularisation in Poisson...	10:00 a.m.
	<i>Coffee break</i>	10:30 a.m.
	Quentin Bertrand Some challenges around retraining generative models on their own data	11:00 a.m.
	Julie Digne Pointwise tools for the analysis of geometric shapes	11:30 a.m.
	<i>Lunch break</i>	12:00 p.m.
	Xiaoqun Zhang Flow based generative models for medical image synthesis	2:00 p.m.
Tuomo Valkonen Differential estimates for fast first-order multilevel nonconvex optimisation	2:30 p.m.	

MONDAY, January 13th, 2025

11:00 a.m.

Silvia Bonettini – Università di Modena e Reggio Emilia,
Dipartimento di Scienze Fisiche, Informatiche e Matematiche
*Linesearch-enhanced forward-backward methods for inexact
nonconvex scenarios*

In the last decades, optimization techniques have been successfully applied to several imaging problems and more and more sophisticated variational models have been proposed in the recent research. In particular, significant improvements of the previous state-of-the-art have been obtained by adopting non-convex settings and combining variational techniques with machine learning approaches. Solving the related optimization problems requires new numerical tools, able to handle their intrinsic difficulties. Indeed, nonconvexity requires specific theoretical and numerical techniques. Moreover, in machine learning settings some elements of the corresponding optimization problem, e.g. the gradient, are often available only as an approximation. In this talk we describe an optimization framework able to handle both nonconvexity and partial knowledge of the variational model. Our approach consists in a forward-backward method with line-search based on approximated values of the objective function and its gradient. As a special case of our general scheme, we derive two algorithms: a line-search based FISTA-like algorithm and a specific inexact method for bilevel optimization problems. The numerical experiments on deblurring and blind deconvolution problems show that the proposed methods are competitive with existing approaches.

11:30 a.m.

Johannes Hertrich – University College London
Importance corrected neural JKO sampling

In order to sample from an unnormalized probability density function, we propose to combine continuous normalizing flows (CNFs) with rejection-resampling steps based on importance weights. We relate the iterative training of CNFs with regularized velocity fields to a proximal mappings in the Wasserstein space. The alternation of local flow steps and non-local rejection-resampling steps allows to overcome local minima and mode collapse for multimodal distributions. The arising model can be trained iteratively, reduces the reverse Kullback-Leibler (KL) loss function in each step, allows to generate iid samples and moreover allows for evaluations of the generated underlying density. Numerical examples demonstrate the efficiency of our approach.

02:30 p.m.

Mame Diarra Fall – Université d'Orléans, Institut Denis Poisson
*Regularization by denoising: Bayesian model
and Langevin-within-split Gibbs sampling*

We propose a Bayesian framework for image inversion by deriving a probabilistic counterpart to the regularization-by-denoising (RED) paradigm. It additionally implements a Monte Carlo algorithm specifically tailored for sampling from the resulting posterior distribution, based on an asymptotically exact data augmentation (AXDA). The proposed algorithm is an approximate instance of split Gibbs sampling (SGS) which embeds one Langevin Monte Carlo step. The proposed method is applied to common imaging tasks such as deblurring, inpainting and super-resolution, demonstrating its efficacy through extensive numerical experiments. References:

- [1] E. C. Faye, M. D. Fall, and N. Dobigeon. "Regularization by denoising: Bayesian model and Langevin-within-split Gibbs sampling". Preprint, 2024.
- [2] S. V. Venkatakrisnan et al. "Plug-and-Play priors for model based reconstruction". In IEEE Global Conf. on Signal and Information Processing, 2013.
- [3] R. Laumont et al. "Bayesian imaging using Plug & Play priors: When Langevin meets Tweedie". SIAM Journal on Imaging Sciences, 2022.
- [4] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)", SIAM Journal on Imaging Sciences, 2017

03:00 p.m.

Marcelo Pereyra – Heriot-Watt University,
School of Mathematical and Computer Sciences
and Maxwell Institute for Mathematical Sciences
*Uncertainty quantification in statistical imaging sciences:
40 years of muddling through*

Probability theory and statistical science are cornerstones of imaging sciences, underpinning many and varied approaches from Markov random fields to score-based denoising diffusion models and stochastic flow-matching techniques. In addition to powerful image estimation methods, statistical science provides a framework for uncertainty quantification and for using image data as quantitative evidence. These capabilities are important for the rigorous interpretation of experimental results and for robust interfacing of quantitative imaging pipelines with scientific and decision-making processes. This talk explores the following question: four decades after the publication of the first seminal papers on the topic, are the probabilities and statistical inferences delivered by existing probabilistic and statistical imaging methods meaningful under replication of an experiment? or are they still only meaningful as subjective measures of belief?

04:00 p.m.

Benedikt Wirth – Universität Münster,
Institute for Applied Mathematics: Analysis and Numerics
The shape space of Sobolev diffeomorphisms and its discretization

A by now classical framework for image registration, computational anatomy and similar applications of mathematical imaging is the large deformation diffeomorphic metric mapping (LDDMM). It is based on equipping a space of diffeomorphisms of chosen regularity with a Riemannian structure that is invariant under composition with diffeomorphisms. The particular setting of Sobolev diffeomorphisms has been quite thoroughly understood during the past decade. This now allows to analyse and devise corresponding numerical discretization schemes and their convergence. I will introduce into the theory of Sobolev LDDMM and present convergence results on its numerics.

04:30 p.m.

Elena di Bernadino – Université Côte d'Azur,
Laboratoire J.A. Dieudonné
*Curvature measures for random excursion sets: theoretical
and computational developments*

The excursion set of a smooth random field carries relevant information in its various geometric measures. Geometric properties of these exceedance regions above a given level provide meaningful theoretical and statistical characterizations for random fields defined on Euclidean domains. Many theoretical results have been obtained for excursions of Gaussian processes and include expected values of the so-called Lipschitz-Killing curvatures (LKC), such as the area, perimeter and Euler characteristic in two-dimensional Euclidean space. In this talk we will describe a recent series of theoretical and computational contributions in this field. Our aim is to provide answers to questions like: i) How the geometric measures of an excursion set can be inferred from a discrete sample of the excursion set; ii) How these measures can be related back to the distributional properties of the random field from which the excursion set was obtained; iii) How the excursion set geometry can be used to infer the extremal behavior of random fields.

TUESDAY, January 14th, 2025

09:00 a.m.

Audrey Repetti – Heriot-Watt University, School of Mathematics and Computer Sciences and School of Engineering and Physical Sciences and Maxwell Institute for Mathematical Sciences
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.

09:30 a.m.

Barbara Pascal – CNRS and Nantes Université, École Centrale Nantes, LS2N
Bilevel optimization for automated data-driven inverse problem resolution

Most inverse problems in signal and image processing are ill-posed. To remove the ambiguity about the solution and design noise-robust estimators, a priori properties, e.g., smoothness or sparsity, can be imposed to the solution through regularization. The main bottleneck to use the obtained variational regularized estimators in practice, i.e., without access to ground truth, is that the quality of the estimates strongly depends on the fine-tuning of the level of regularization. A classical approach to automated and data-driven selection of regularization parameter consists in designing a data-dependent unbiased estimator of the error, the minimization of which provides an approximate of the optimal parameters. The resulting overall procedure can be formulated as a bilevel optimization problem, the inner loop computing the variational regularized estimator and the outer loop selecting hyperparameters. The design of a fully automated data-driven procedure adapted to inverse problems corrupted with highly correlated noise will be described in detail and exemplified on a texture segmentation problem. Its applicability to other inverse problems will be demonstrated through numerical simulations on both synthetic and real-world data.

10:00 a.m.

Elisa Riccietti – ENS Lyon, LIP
Exploiting multiple resolutions to accelerate inverse problems in imaging

Solving large scale optimization problems is a challenging task and alleviating its computational cost is an open research problem. In this talk we propose a method to face this challenge, which is based on an idea that is at the core of multilevel optimization methods: exploiting the structure of the problem to define coarse approximations of the objective function, representing the problem at different resolutions. We present IML FISTA, a multilevel inertial proximal algorithm to tackle non-smooth problems that draws ideas from the multilevel setting for smooth optimization. IML FISTA is able to handle state-of-the-art regularization techniques such as total variation and non-local total-variation, while providing a relatively simple construction of coarse approximations. We demonstrate the effectiveness of the approach on color and hyperspectral images reconstruction problems.

11:00 a.m. | **Kimia Nadjahi** – ENS Paris, DI
Scalable unbalanced optimal transport by slicing

Substantial advances have been made in designing optimal transport (OT) variants which are either computationally and statistically more efficient or robust. Among them, sliced OT distances have been extensively used to mitigate the cubic algorithmic complexity and curse of dimensionality of OT. In parallel, unbalanced OT was designed to allow comparisons of more general positive measures, while being robust to outliers. In this talk, we bridge the gap between those two concepts and develop a general framework for efficiently comparing positive measures. We formulate two different versions of "sliced unbalanced OT" and study the associated topology and statistical properties. We then develop a GPU-friendly Frank-Wolfe algorithm to compute the two corresponding loss functions. The resulting methodology is modular, as it encompasses and extends prior related work, and brings the cubical computational complexity down to almost linear $O(n \log(n))$ when comparing two discrete measures supported on n points. We finally conduct an empirical analysis of our methodology on both synthetic and real datasets, to illustrate its computational efficiency, relevance and applicability to real-world scenarios, including color transfer and barycenters of geophysical data.

Reference: "Slicing Unbalanced Optimal Transport", Clément Bonet, Kimia Nadjahi, Thibault Séjourné, Kilian Fatras, Nicolas Courty (2024) (<https://arxiv.org/abs/2306.07176>)

11:30 a.m. | **Matthew Thorpe** – University of Warwick, Department of Statistics
How many labels do you need in semi-supervised learning?

Semi-supervised learning (SSL) is the problem of finding missing labels from a partially labelled data set. The heuristic one uses is that "similar feature vectors should have similar labels". The notion of similarity between feature vectors explored in this talk comes from a graph-based geometry where an edge is placed between feature vectors that are closer than some connectivity radius. A natural variational solution to the SSL is to minimise a Dirichlet energy built from the graph topology. And a natural question is to ask what happens as the number of feature vectors goes to infinity? In this talk I will give results on the asymptotics of graph-based SSL using an optimal transport topology. The results will include a lower bound on the number of labels needed for consistency.

02:00 p.m. | **Antonin Chambolle** – CNRS and Université Paris-Dauphine-PSL, CEREMADE
Some properties of the solutions of Total-Variation regularized inverse problems

The total variation has been successful as a regularizer for inverse problems in imaging, thanks to its ability to preserve discontinuities (edges) and its relative simplicity (convexity). Even if largely outdated by deep learning based methods, it still can be useful in some regimes (low noise, large scale images). This talk is about the preservation of edges in total-variation based denoising. We revisit old proofs which show in some settings that no spurious edges are created by this approach. Our new approach, based on works of T. Valkonen, is natural and simple and applies to more settings (color/multispectral data, some higher order models, non-local variants). This is joint work with Michał Łasica from Warsaw and Konstantinos Bessas (Pavia).

02:30 p.m. | **Clarice Poon** – University of Warwick, Mathematics Institute
Hadamard Langevin dynamics for sampling sparse priors

Priors with non-smooth log densities have been widely used in Bayesian inverse problems, particularly in imaging, due to their sparsity inducing properties. To date, the majority of algorithms for handling such densities are based on proximal Langevin dynamics where one replaces the non-smooth part by a smooth approximation known as the Moreau envelope. In this work, we introduce a novel approach for sampling densities with ℓ_1 -priors based on a Hadamard product parameterization. This builds upon the idea that the Laplace prior has a Gaussian mixture representation and our method can be seen as a form of overparametrization: by increasing the number of variables, we construct a density from which one can directly recover the original density. This is fundamentally different from proximal-type approaches since our resolution is exact, while proximal-based methods introduce additional bias due to the Moreau-envelope smoothing. For our new density, we present its Langevin dynamics in continuous time and establish well-posedness and geometric ergodicity. We also present a discretization scheme for the continuous dynamics and prove convergence as the time-step diminishes.

03:00 p.m. | **Matthias J. Ehrhardt** – University of Bath,
Department of Mathematical Sciences
Inexact algorithms for bilevel learning

Variational regularization techniques are dominant in the field of inverse problems. A drawback of these techniques is that they are dependent on a number of parameters which have to be set by the user. This issue can be approached by machine learning where we estimate these parameters from data. This is known as "Bilevel Learning" and has been successfully applied to many tasks, some as small-dimensional as learning a regularization parameter, others as high-dimensional as learning parameters of neural networks. While mathematically appealing, this strategy leads to a nested optimization problem which is practically challenging since function values and gradients cannot be computed to a high-enough precision. In this talk we discuss new computational approaches for this problem which do not assume exact knowledge of these. It turns out that a clever choice on the accuracy leads to much faster yet stable and robust solutions.

WEDNESDAY, January 15th, 2025

09:00 a.m. | **Jonathan Dong** – EPFL, Biomedical Imaging Group
Random phase retrieval: theory and implementation

Phase retrieval consists in the recovery of a complex-valued signal from intensity-only measurements. Strong results have recently been obtained for random models with i.i.d. sensing matrix components. In this presentation, we will review phase retrieval under a unifying framework for its different applications and focus on the random case. We will describe the latest theoretical results and limitations in practice, to conclude with structured random models combining the efficiency of fast Fourier transforms and the robustness of random phase retrieval reconstructions.

09:30 a.m. | **Anne Wald** – University of Göttingen,
Institute for Numerical and Applied Mathematics
*Nano-CT imaging as an inverse problem
with inexact forward operator*

Tomographic X-ray imaging on the nano-scale is an important tool to visualize the structure of materials such as alloys or biological tissue. Due to the small scale on which the data acquisition takes place, small perturbances caused by the environment become significant and cause a motion of the object relative to the scanner during the scan. Since this motion is hard to estimate and its incorporation into the reconstruction process strongly increases the numerical effort, we aim at a different approach for a stable reconstruction: We interpret the object motion as a modelling inexactness in comparison to the model in the static case and aim at compensating this modelling error. Apart from an unknown object motion, there are further properties of the physical setup that cause a deviation from the standard Radon-type forward operator. We discuss recent advances in compensating different kinds of modelling errors in nano-CT and present reconstructions from measurement data.

10:00 a.m. | **Luca Calatroni** – CNRS and Université Côte d'Azur, Laboratoire I3S
*Deep equilibrium mirror descent for learning image regularisation
in Poisson inverse problems*

We consider the framework of deep equilibrium models to learn image regularisers in the context of imaging data corrupted by Poisson noise. In this framework, the Kullback-Leibler divergence is usually considered as a data term, which poses some difficulties due to its lack of Lipschitz-smoothness around zero. By using a mirror descent algorithm and enforcing a Lipschitz-like condition to guarantee convergence of the scheme even in non L-smooth settings, we propose a training strategy that learns efficiently the regulariser using limited data and reduced computational times. The key ingredients of the proposed approach are an efficient forward step obtained by means of an efficient backtracking strategy and a cheap backward step relying on Jacobian-free approximations. Numerical results on exemplar image denoising/deblurring problems and some open questions are presented. This is joint work with C. Daniele, S. Villa (MaLGA, University of Genoa, IT) and S. Vaiter (LJAD, CNRS, FR).

11:00 a.m. | **Quentin Bertrand** – Inria
Some challenges around retraining generative models on their own data

Deep generative models have made tremendous progress in modeling complex data, often exhibiting generation quality that surpasses a typical human's ability to discern the authenticity of samples. Undeniably, a key driver of this success is enabled by the massive amounts of web-scale data consumed by these models. Due to these models' striking performance and ease of availability, the web will inevitably be increasingly populated with synthetic content. Such a fact directly implies that future iterations of generative models will be trained on both clean and artificially generated data from past models. In addition, in practice, synthetic data is often subject to human feedback and curated by users before being used and uploaded online. For instance, many interfaces of popular text-to-image generative models, such as Stable Diffusion or Midjourney, produce several variations of an image for a given query which can eventually be curated by the users. In this talk we will discuss the impact of training generative models on mixed datasets—from classical training on real data to self-consuming generative models trained on purely synthetic curated data.

11:30 a.m. | **Julie Digne** – CNRS and Université Lyon 1, LIRIS
Pointwise tools for the analysis of geometric shapes

Digitized geometric shapes are omnipresent in issues of heritage preservation, or for the digital creation of industrial parts. They are the result of acquisition processes producing a set of 3d points possibly noisy or only partially covering the shape. In this talk, I will present new tools for the analysis of shapes allowing us to highlight interesting local differential and frequency properties. This was previously used as a way to exaggerate details of a shape. In this talk, I will present how this framework allows for an elegant formulation of shape characteristic lines.

02:00 p.m. | **Xiaoqun Zhang** – Shanghai Jiao Tong University
Flow based generative models for medical image synthesis

The synthesis of high-quality medical images is critical for enhancing clinical decision-making, diagnostic accuracy, and treatment planning, as well as for applications such as data augmentation and image quality improvement. Flow-based generative models have demonstrated significant potential in modeling complex data distributions and generating realistic synthetic images. This talk presents two novel approaches that contribute to advancements in flow-based generative modeling for medical image synthesis. The first approach introduces SyMOT-Flow, an invertible transformation model that minimizes the symmetric maximum mean discrepancy between samples from two unknown distributions, incorporating an optimal transport cost as regularization. This ensures short-distance and interpretable mappings, leading to more stable and accurate sample generation. The model is validated through low-dimensional illustrative examples and high-dimensional bi-modality medical image generation tasks. The second approach proposes Bi-DPM (Bi-directional Discrete Process Matching), a novel model for bi-modality image synthesis. Unlike traditional flow-based methods that rely on computationally intensive ordinary differential equation (ODE) solvers, Bi-DPM utilizes forward and backward flows with enhanced consistency over discrete time steps. This results in efficient and high-quality image synthesis guided by paired data. Experimental results on MRI T1/T2 and CT/MRI datasets show that Bi-DPM achieves superior image quality and accurately synthesizes anatomical regions compared to existing methods. These contributions offer practical advancements in flow-based medical image synthesis, addressing computational efficiency and image fidelity while providing tools that can support improved clinical workflows and outcomes.

02:30 p.m.

Tuomo Valkonen – Escuela Politécnica Nacional, ModeMat and University of Helsinki, Department of Mathematics and Statistics
Differential estimates for fast first-order multilevel nonconvex optimisation

PDE constraints appear in inverse imaging problems as physical models for measurements, while bilevel optimisation can be used for optimal experimental design and parameter learning. Such problems have been traditionally very expensive to solve, but recently, effective single-loop approaches have been introduced, both in our work, as well as in the machine learning community. In this talk, we discuss a simple gradient estimation formalisation for very general single-loop methods that include primal-dual methods for the inner problem, and conventional iterative solvers (Jacobi, Gauss–Seidel, conjugate gradients) for the adjoint problem and PDE constraints. This talk is based on joint work with Neil Dizon, Bjørn Jensen, and Ensio Suonperä.

The Biennial French-German Mathematics in Imaging Ph.D. Prize is organized by CNRS RT MAIAGES, GAMM-SIP, and SMAI SIGMA group. It is sponsored by RT MAIAGES and Fraunhofer ITWM. It is awarded to a young researcher with outstanding contribution to the field of Mathematics in Imaging.

Jury Members

Kristian Bredies
– University of Graz
Serena Morigi
– University of Bologna
Marcelo Pereyra
– Heriot-Watt University
Xiaoqun Zhang
– Shanghai Jiao Tong University

Former winners

2023: Sebastian Jonas Neumayer
2021: Leon Bungert
2018: Emmanuel Soubiès



We are pleased to announce that Florian Beier has been awarded the 2025 MIA PhD prize!

Florian Beier received the Ph.D. degree in mathematics under supervision of Gabriele Steidl from the Technische Universität Berlin, Germany. He is interested in topics related to computer vision, dynamical systems and statistical inference. His contributions are based on the analysis of generalized optimal transport spaces and their differential geometry.

He is giving a talk on **January 13th at 2 p.m.**

A Geometric Optimal Transport Framework for 3D Shape Interpolation

The Gromov-Wasserstein (GW) transport problem is a generalization of the classic optimal transport problem, which seeks a relaxed correspondence between two measures while preserving their internal geometry. Due to meeting this theoretical underpinning, it is a valuable tool for the analysis of objects that do not possess a natural embedding or should be studied independently of it. Prime applications can thus be found in e.g. shape matching, classification and interpolation tasks. To tackle the latter, one theoretically justified approach is the employment of GW barycenters, which are generalized Fréchet means with respect to the GW distance.

After giving a gentle and illustrative introduction to the GW transport problem, we turn our attention to GW barycenters. Motivated by obtaining a numerically tractable method for their computation, we study the geometry of the induced GW space. Our theoretical results in this context allow us to lift a known fixpoint iteration for the computation of Fréchet means in Riemannian manifolds to the GW setting. The lifted iteration is simple to implement in practice and monotonically improves the quality of the barycenter. We provide numerical evidence of the potential of this method, including multi 3d shape interpolations.

TUESDAY, January 14th, 2025

Carl Allen – ENS Paris, DI

Unpicking data at the seams: vaes, disentanglement and independent components

Disentanglement, or identifying salient statistically independent factors of the data, is of interest in many areas of machine learning and statistics, such as synthetic data generation with controlled properties, robust classification of features, parsimonious encoding, and improving our understanding of the generative process underlying the data. Disentanglement is observed in several generative paradigms, including Variational Autoencoders (VAEs), Generative Adversarial Networks and diffusion models. Particular progress has recently been made in understanding disentanglement in VAEs, where the choice of diagonal posterior covariance matrices is proposed to promote mutual orthogonality between columns of the decoder's Jacobian. We continue this thread to show how such linear independence translates to statistical independence, completing the chain in understanding how the VAE's objective identifies independent components of, or disentangles, the data.

Jonathan Brokman – Technion - Israel Institute of Technology

Manifold induced biases for zero-shot detection of generated images

Diffusion models, the leading image generative models of today, estimate the score function—the gradient of the log probability of data—solely through optimization on data samples, without requiring access to the true underlying probability distribution. This work investigates whether this estimated score function can be used to compute higher-order differentials, such as the total-variation sub-gradient, to capture biases in the learned probability manifold under zero-shot settings with a pre-trained diffusion model. Similar zero-shot analyses has previously proven effective for detecting generated text, and here we extend these principles to diffusion models and the image domain, combining novel theoretical derivations with practical estimations. We validate our approach through experiments comparing analytical expressions, exploring structured cases like Gaussian Mixture Models, and conducting ablation studies on key hyperparameters (e.g., sample size, local neighborhood size). Lastly, we demonstrate its applicability to natural image data, underscoring its potential for analyzing practical diffusion models and detecting generated images.

Nathan Buskulic – Université de Caen, GREYC

Guarantees of self-supervised networks for inverse problems trained with a time-continuous inertial system

In this work we study the optimization trajectory of the parameters of a network when one uses a time-continuous inertial dynamical system to train a neural network. More specifically, we will focus on an inertial system combining viscous and Hessian-driven damping in the case where the network is trained to minimize a mean square error loss function to solve an inverse problem. We show that for the right choice of parameters for the inertia and Hessian damping, one can obtain accelerated convergence, hence recovery, guarantees compared to gradient flow, at the cost of a stronger initialization condition. We also demonstrate how the overparametrization bound for a two-layer Deep Inverse Prior is affected by this stronger initialization requirement.

Nathanaël Cuvelle–Magar – ENS Paris, DI
*Optimal denoising of geometrically regular images
 with scattering coefficients*

Optimal suppression of additive Gaussian white noise has many image processing applications and is a key step to generate images with score diffusion algorithms. For piecewise regular images with edges, nearly optimal denoisers can be computed by thresholding a sparse representation in a dictionary of curvelets or bandlets. It requires to adapt the support of selected dictionary vectors to geometric image properties. In contrast, convolutional deep neural networks can implement optimal denoising algorithms with a cascade of fixed support filters and non-linearities. It also applies to much more complex images. An outstanding issue is to understand how convolutional neural networks can reach optimal results, without adapting their filter support. We show that a two layers network with wavelet filters can disentangle geometric support constraints and image regularity properties. This scattering transform explains why filter supports do not need to be adapted to the image geometry. For piecewise C^α images having piecewise C^α edges, numerical results are asymptotically optimal. This is supported by a characterization of the scattering coefficient decay, but this denoising optimality remains a mathematical conjecture.

Jay Dhesi – University of Cambridge
*Relational persistent homology for multispecies data
 with application to the tumor microenvironment*

Topological data analysis (TDA) is an active field of mathematics for quantifying shape in complex data. Standard methods in TDA such as persistent homology (PH) are typically focused on the analysis of data consisting of a single entity (e.g., cells or molecular species). However, state-of-the-art data collection techniques now generate exquisitely detailed multispecies data, prompting a need for methods that can examine and quantify the relations among them. Such heterogeneous data types arise in many contexts, ranging from biomedical imaging, geospatial analysis, to species ecology. Here, we propose two methods for encoding spatial relations among different data types that are based on Dowker complexes and Witness complexes. We apply the methods to synthetic multispecies data of a tumor microenvironment and analyze topological features that capture relations between different cell types, e.g., blood vessels, macrophages, tumor cells, and necrotic cells. We demonstrate that relational topological features can extract biological insight, including the dominant immune cell phenotype (an important predictor of patient prognosis) and the parameter regimes of a data-generating model. The methods provide a quantitative perspective on the relational analysis of multispecies spatial data, overcome the limits of traditional PH, and are readily computable.

Yara Elshiaty – Heidelberg University, Institute of Mathematics
Multilevel mirror descent

We present the Multilevel Mirror Descent (MLMD) method, a novel multilevel optimization framework tailored to constrained convex problems with relative Lipschitz smoothness. Our approach extends the classical multilevel optimization framework (MGOPT) to handle Bregman-based geometries and constrained domains. We provide a rigorous analysis of MLMD for multiple coarse levels and establish a global sublinear convergence rate. We demonstrate the effectiveness of MLMD in the context of image reconstruction, providing theoretical guarantees for the well-posedness of the multilevel framework and validating its performance through numerical experiments. Additional examples beyond imaging further showcase its versatility.

Anne Gagneux – ENS Lyon
Convexity in ReLU Neural Networks: beyond ICNNs?

Convex functions and their gradients play a critical role in mathematical imaging, from proximal optimization to Optimal Transport. The successes of deep learning has led many to use learning based methods, where fixed functions or operators are replaced by learned neural networks. Regardless of their empirical superiority, establishing rigorous guarantees for these methods often requires to impose structural constraints on neural architectures, in particular convexity. The most popular way to do so is to use so-called Input Convex Neural Networks (ICNNs). In order to explore the expressivity of ICNNs, we provide necessary and sufficient conditions for a ReLU neural network to be convex. Such characterizations are based on product of weights and activation, and write nicely for any architecture in the path-lifting framework. As particular applications, we study our characterizations in depth for 1 and 2-hidden-layer neural network: we show that every convex function implemented by a 1 hidden layer ReLU network can be also expressed by an ICNN with the same architecture; but this property no longer holds with more layers. Finally, we provide a numerical procedure that allows an exact check of convexity for ReLU neural networks with a large number of affine regions.

Nicolas Gindrier – Austrian Academy of Sciences, RICAM
*Exact reconstruction for saddle trajectory
with transverse truncation in CBCT*

In cone beam (CB) tomography, the saddle trajectory is a X-ray source trajectory quite studied, currently undergoing industrial development. It can be described as a sine on a cylinder. However, the literature focuses more on the framework of no truncation or axial truncation. Truncation occur when some lines intersecting the object and the X-ray source are not measured. The axial truncation come for example when the detector height is too small. For material constraints or dose reduction, transverse truncation can appear. In this case, the most relevant and suitable analytical method is the differential backprojection (DBP) method, with its extension called M-lines method. We proved the DBP method is well-suited for saddle trajectories with transverse truncation. The reconstruction can be done in any region of interest being inside the intersection of the field of view and the convex hull of the saddle trajectory.

Sékolène Martin – TU Berlin
*PnP-Flow: Plug-and-Play image restoration
with flow matching*

In this paper, we introduce Plug-and-Play (PnP) Flow Matching, an algorithm for solving imaging inverse problems. PnP methods leverage the strength of pre-trained denoisers, often deep neural networks, by integrating them in optimization schemes. While they achieve state-of-the-art performance on various inverse problems in imaging, PnP approaches face inherent limitations on more generative tasks like inpainting. On the other hand, generative models such as Flow Matching pushed the boundary in image sampling yet lack a clear method for efficient use in image restoration. We propose to combine the PnP framework with Flow Matching (FM) by defining a time-dependent denoiser using a pre-trained FM model. Our algorithm alternates between gradient descent steps on the data-fidelity term, rejections onto the learned FM path, and denoising. Notably, our method is computationally efficient and memory-friendly, as it avoids backpropagation through ODEs and trace computations. We evaluate its performance on denoising, super-resolution, deblurring, and inpainting tasks, demonstrating superior results compared to existing PnP algorithms and Flow Matching based state-of-the-art methods.

Alex Massucco – University of Cambridge
*Manifold-induced properties for modified
ResNet-like neural networks*

In recent literature, great mathematical effort has been made to deepen our knowledge and understanding of the theoretical behavior of neural networks. In particular, structure-preserving architectures, trying to encode or force on the network architecture some additional *a priori* problem-dependent properties, have become increasingly popular. In this work, we investigate a new family of ResNet-like architectures based on data-learned manifolds. The main advantage of this method is represented by the possibility of manipulating the forward pass, from inputs to outputs, typical of ResNet architectures incorporating the additional structure on the manifold without modifying the underlying vector field inducing the flow. This permits interpreting the learning process of the additional prescribed properties as the action of a homotopy map in the space of neural networks only deforming the underlying geometric substrates and preserving the learned task-related weights. We provide an explicit method to construct these architectures and present a training algorithm to achieve a satisfactory balance between the performances and the prescribed properties. We focus in particular on symplectic networks and networks less sensitive to input perturbation. On the one hand, we consider experiments demonstrating an improvement in the long-time behavior for networks approximating Hamiltonian dynamics trained on short-time trajectories. On the other hand, we provide robustness improvements in the context of networks applied to noisy data on some common benchmarks.

Mimoun Mohamed – Université Aix-Marseille,
Laboratoire d'Informatique et Systèmes
& Institut de Mathématiques de Marseille
*Straight-through meets sparse recovery:
the support exploration algorithm*

The straight-through estimator (STE) is commonly used to optimize quantized neural networks, yet its contexts of effective performance are still unclear despite empirical successes. To make a step forward in this comprehension, we apply STE to a well-understood problem: sparse support recovery. We introduce the Support Exploration Algorithm (SEA), a novel algorithm promoting sparsity, and we analyze its performance in support recovery (a.k.a. model selection) problems. SEA explores more supports than the state-of-the-art, leading to superior performance in experiments, especially when the columns of A are strongly coherent. The theoretical analysis considers recovery guarantees when the linear measurements matrix A satisfies the Restricted Isometry Property (RIP). The sufficient conditions of recovery are comparable but more stringent than those of the state-of-the-art in sparse support recovery. Their significance lies mainly in their applicability to an instance of the STE.

Saskia Neuber – University of Lübeck,
Institute of Mathematics and Image Computing
*A Multigrid approach for a coupled second-order
regularizer in image registration*

Image registration is needed in a various number of fields e.g. in medical imaging. The goal of image registration is to determine a deformation between two corresponding images. We follow a variational approach, where a joint functional is defined, consisting of a data similarity term and a regularizer. Here we focus on a family of vector field (VF) regularizers, based on second-order derivatives. Second-order regularization naturally leads to smoother solutions as compared to first-order and enables a proper treatment of constraints such as landmark matches or 2D/3D registration. The price for the second order is that the numerical treatment of the registration problem is significantly more complex. In particular the VF regularizers consist of the gradient of the divergence and of the rotation of the deformation. These regularizers enable a coupling of the components and include as a special cases curvature regularization. A typical optimization strategy for the minimization of the joint functional in image registration are Newton-type algorithms. Here the most costly part is solving the potentially large linear system, which can be regarded as a discretization of a linearization of a partial differential equation of fourth order. Multigrid methods are known for its high efficiency in determining a solution of a system of linear equations. Hence we propose an efficient multigrid solver for the image registration problem containing a VF regularizer. In particular, we present a local Fourier analysis to show that the suggested discretization of the VF regularizer is well suited for multigrid methods. More specifically, we provide an explicit number for the h -ellipticity measure and the smoothing factor for a collective ω -relaxed Jacobi-type iteration. Our numerical results, including 3D image registration tasks, underline that the MG solver has in fact a complexity of $\mathcal{O}(n)$, where n is the number of unknowns.

Émile Pierret – Université d'Orléans, Institut Denis Poisson
*Diffusion models for Gaussian distributions:
exact solutions and Wassersteins errors*

Diffusion or score-based models recently showed high performance in image generation. They rely on a forward and a backward stochastic differential equations (SDE). The sampling of a data distribution is achieved by solving numerically the backward SDE or its associated flow ODE. Studying the convergence of these models necessitates to control four different types of error: the initialization error, the truncation error, the discretization and the score approximation. In this paper, we study theoretically the behavior of diffusion models and their numerical implementation when the data distribution is Gaussian. In this restricted framework where the score function is a linear operator, we derive the analytical solutions of the backward SDE and the probability flow ODE. We prove that these solutions and their discretizations are all Gaussian processes, which allows us to compute exact Wasserstein errors induced by each error type for any sampling scheme. Monitoring convergence directly in the data space instead of relying on inception features, our experiments show that the recommended numerical schemes from the diffusion models literature are also the best sampling schemes for Gaussian distributions.

Marien Renaud – Université de Bordeaux,
Institut de Mathématiques de Bordeaux
*Convergence analysis of a proximal stochastic
denoising regularization algorithm*

Plug-and-Play methods for image restoration are iterative algorithms that solve a variational problem to restore an image. These algorithms are known to be flexible to changes of degradation and to perform state-of-the-art restoration. Recently, a lot of efforts have been made to explore new algorithms to solve this variational problem based on the Plug-and-Play or REgularization by Denoising (RED) frameworks, such as SNORE that is a converging stochastic gradient descent algorithm. A variant of this algorithm, named SNORE Prox, reaches state-of-the-art performances, especially for inpainting tasks. However, the convergence of SNORE Prox, that can be seen as a stochastic proximal gradient descent, has not been analyzed so far. In this paper, we prove the convergence of SNORE Prox under reasonable non-convex assumptions.

Pia Schulz – University of Lübeck,
Institute of Mathematics and Image Computing
Appropriate Regularization Order in 3D/2D Image Registration

We investigate the impact of the regularization order in variational 3D/2D image registration. Such registration problems often occur in image-guided interventions, e.g. in cardiology, neurology, or orthopedics. Usually, 3D images are acquired pre-operatively in order to plan the treatment, but during the intervention only 2D images are available, e.g. due to time or dose constraints. Obviously, using the 3D data in combination with the intra-operative 2D data allows for easier guidance and monitoring during the surgery. 3D/2D image registration enables this by aligning the 3D data to the 2D data. A common approach to calculate the 3D deformation is to minimize an energy that consists of a data fitting and a regularization term. While the first is evaluated on a 2D set, the regularization acts on a 3D domain. Thus, the regularizer has to ensure that the information given on the 2D set is transferred to the 3D domain. Therefore, the regularization should be chosen such that the resulting 3D deformation is sufficiently smooth, e.g. continuously differentiable. In this work, we address the question which regularization order is required to guarantee continuously differentiable deformations. Our main result is that specific second-order regularizers are sufficient, while first-order regularizers are generally not.

Bernardin Tamo Amougou – Heriot-Watt University
Self-supervised conformal prediction for uncertainty quantification in imaging problems

Most image restoration problems are ill-conditioned or ill-posed and hence involve significant uncertainty. Quantifying this uncertainty is crucial for reliably interpreting experimental results, particularly when reconstructed images inform critical decisions and science. However, most existing image restoration methods either fail to quantify uncertainty or provide estimates that are highly inaccurate. Conformal prediction has recently emerged as a flexible framework to equip any estimator with uncertainty quantification capabilities that, by construction, have nearly exact marginal coverage. To achieve this, conformal prediction relies on abundant ground truth data for calibration. However, in image restoration problems, reliable ground truth data is often expensive or not possible to acquire. Also, reliance on ground truth data can introduce large biases in situations of distribution shift between calibration and deployment. This paper seeks to develop a more robust approach to conformal prediction for image restoration problems by proposing a self-supervised conformal prediction method that leverages Stein's Unbiased Risk Estimator (SURE) to self-calibrate itself directly from the observed noisy measurements, bypassing the need for ground truth. The method is suitable for any linear imaging inverse problem that is ill-conditioned, and it is especially powerful when used with modern self-supervised image restoration techniques that can also be trained directly from measurement data. The proposed approach is demonstrated through numerical experiments on image denoising and deblurring, where it delivers results that are remarkably accurate and comparable to those obtained by supervised conformal prediction with ground truth data.

Hong Ye Tan – University of Cambridge
Unsupervised training of convex regularizers using maximum likelihood estimation

Unsupervised learning is a training approach in the situation where ground truth data is unavailable, such as inverse imaging problems. We present an unsupervised Bayesian training approach to learning convex neural network regularizers using a fixed noisy dataset, based on a dual Markov chain estimation method. Compared to classical supervised adversarial regularization methods, where there is access to both clean images as well as unlimited to noisy copies, we demonstrate close performance on natural image Gaussian deconvolution and Poisson denoising tasks.

Corentin Vazia – Université Bretagne-Sud,
Laboratoire de Mathématiques de Bretagne Atlantique
Guidance of a diffusion model for material decomposition in photon-counting computed tomography

Computed tomography image reconstruction is a well-known inverse problem in medical imaging that aims to retrieve the linear attenuation map of the scanned object or patient. In the context of spectral computed tomography, and more specifically with photon-counting detectors, we can now obtain energy-dependent linear attenuation coefficients (LAC) from transmission measurements at different energy levels. This (already ill-posed) inverse problem can be followed by material decomposition where we decompose the LAC images into a material attenuation basis. Recent advances in diffusion models can be used as priors and for regularization of inverse problems by guiding the sampling process (e.g. with diffusion posterior sampling) to a solution that fits with the measurements. In this work, we present and compare multiple methods using diffusion posterior sampling for material decomposition.

The conference will take place at the




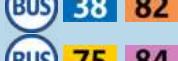


Institute Henri Poincaré (IHP)
 Borel building, Hermite lecture hall
 11 Rue Pierre et Marie Curie
 75005 Paris

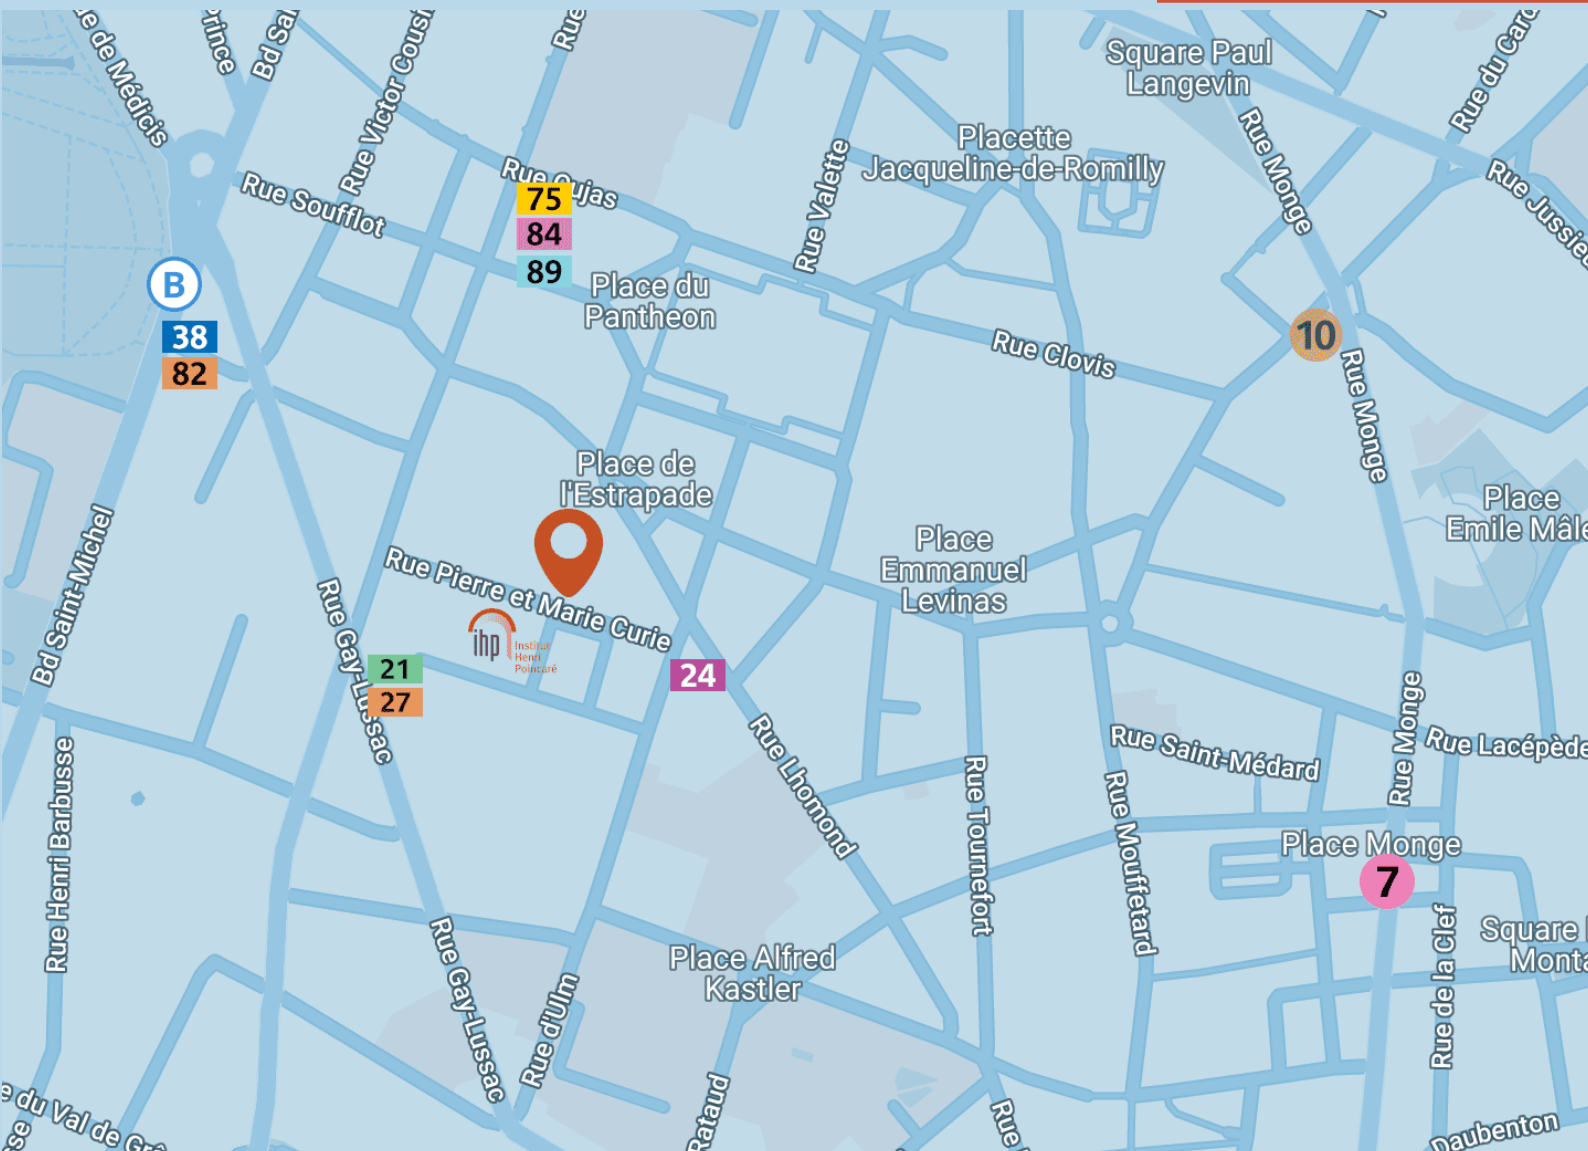
located in the Latin Quarter of Paris, close to the Panthéon.

ACCESS





To go to the Institute Henri Poincaré, you can take :

-  Luxembourg
-  Cardinal Lemoine
-  Place Monge
-  Saint-Jacques - Gay Lussac
-  Musée et Institut Curie
-  Luxembourg
-  Panthéon

Note: Since January 1st, the public transport pricing has changed: there are **two** kind of tickets, one for rail network (metro/RER) and one for bus/tram network. Note that all tickets can be loaded on a smart phone or a Navigo Easy pass (€2), but different types of tickets cannot be loaded on the same medium the same time.
 More information on: <https://www.ratp.fr/en>



KEY

ICON	TYPE	NOTES
	ATMs	
	Grocery stores	
	Pharmacies	Closest ones: Pharmacie Gay Lussac, 38 rue Gay-Lussac Pharmacie du Panthéon, 171 rue Saint-Jacques Pharmacie Saint Jacques, 243 rue Saint-Jacques
	Restaurants & bakeries*	Vegetarian**: Garden Café Paris, 220 rue Saint-Jacques Gluten-free**: Kim Ly, 15 rue Royer-Collard

* This is a list of nearby restaurants and bakeries for your convenience. Please note that these are indicative options and not specific recommendations.

** According to Google

