Generative modeling for constrained Bayesian inverse problems in imaging

PhD Thesis, Start Spring 2026

IMT Atlantique, Brest Campus, SPIN Doctoral School, Funding: Cluster SequoIA

Keywords Inverse problems, Generative models, Flow matching, Uncertainty quantification, Differential Geometry, Remote Sensing, Spectral Unmixing.

Academic Context. IMT Atlantique, internationally recognized for the quality of its research, is a leading general engineering school under the aegis of the Ministry of Industry and Digital Technology, ranked in the three main international rankings (THE, SHANGHAI, QS). Located on three campuses, Brest, Nantes and Rennes, IMT Atlantique aims to combine digital technology and energy to transform society and industry through training, research and innovation. It aims to be the leading French higher education and research institution in this field on an international scale. With 290 researchers and permanent lecturers, 1000 publications and 18 M€ of contracts, it supervises 2300 students each year and its training courses are based on cutting-edge research carried out within 6 joint research units: GEPEA, IRISA, LATIM, LABSTICC, LS2N and SUBATECH. The proposed thesis is part of the research activities of the team ODYSSEY and of the laboratory Lab-STICC and the department Mathematical and Electrical Engineering (MEE).

Applicative Context. An inverse problem aims at estimating model parameters from input data, having access to a model describing how to generate the observations if the parameters to estimate were known. For instance, in optical remote sensing (e.g. hyperspectral imagery), the unmixing problem aims at separating the contributions of the different materials that are present in the field of view of each pixel, by estimating the signatures of the different materials and their relative proportions in every pixel of the image (Fig. 1 and [1]). Applications include environmental monitoring, land cover estimation, physical parameter estimation, among others. A typical observation model (giving the likelihood of the data $p(\mathbf{X}|\mathbf{A})$) states that the different materials contribute linearly in each pixel:

$$\mathbf{X} = \mathbf{S}\mathbf{A} + \mathbf{E}$$

where $\mathbf{X}$ gathers all the L-dimensional pixels (L is the number of channels) of the image in a matrix, $\mathbf{S}$ is a matrix containing the signatures of the materials that are present in the image, $\mathbf{A}$ gathers the proportions of every material in every pixel, and $\mathbf{E}$ is a spatially and spectrally white Gaussian noise. This problem can be ill-posed, in particular for low spatial resolutions, high levels of noise or partial observations, so a prior distribution $p(\mathbf{A})$ is typically incorporated in a Bayesian framework (e.g. enforcing that neighboring pixels are highly correlated). An additional difficulty here is that $\mathbf{A}$ is a structured geometric object: an image for which each pixel is constrained to belong to the probability simplex. In the absence of reliable ground truth, an important feature of a solver is to be able to quantify the uncertainty in the estimates of $\mathbf{A}$, by sampling the posterior distribution $p(\mathbf{A}|\mathbf{X})$.

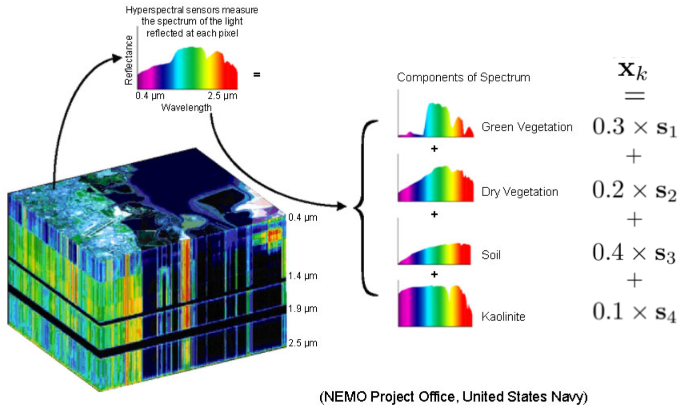

Figure 1: Hyperspectral image unmixing: each pixel $\mathbf{x}_k$, $k = 1, ..., N$ is decomposed into a linear combination of reference signatures, with coefficients that are positive and sum to one.

Methodological and technical challenges. The objectives of the thesis are threefold:

- Design priors on the proportions $\mathbf{A}$ that are as realistic as possible and can handle the constraints.
- Integrate them into Bayesian models to sample the posterior distribution and provide uncertainty quantification on the estimated parameters.
- Generalize and extend the constructions to other types of geometric constraints and applications.

We have identified a way to define unsupervised suitable prior distributions on $\mathbf{A}$ by adapting Gaussian Random Fields [2], which is a first step towards the objective. However, in spite of many favorable properties, such unsupervised priors may not be representative of real world spatialized distributions for the proportions. Thus, following a recent approaches [3], we aim to define supervised priors through modern neural-based generative models, in particular flow matching [4, 5]. These priors are able to generate samples matching the distribution of a training dataset. Some recent breakthroughs handle distributions on structured supports [6, 7], but these models are not suited to generate 2D fields (images) that are subject to constraints. The goal of this thesis is to design such models and to mobilize them for the unmixing problem. Depending on the findings and the candidate's interests, other types of geometric constraints and applications can be envisioned, ranging from oceanography, medical imaging, or uncertainty quantification in general.

Work plan

- Months 1-6: Study of generative models for images (Gaussian Processes, Diffusion models, Flow Matching) and their applications to Bayesian inverse problems, and Literature review around constrained generative modeling or sampling/inference. Usage of GP models as an unsupervised baseline.
- Months 6-18: Development of Generative models for simplex-valued images and integration into inverse problems pipelines. Data preparation for supervised model training. Application to the unmixing problem. First valorization of the methods via submissions to top journals and conferences in AI/ML or image processing.
- Months 18-30: Investigation of various applications and other types of constraints. Submissions to top journals and conferences in AI/ML or applicative domains.
- Months 30-36: Thesis manuscript and defense preparation.

Profile of the candidate. We expect the candidate (final year engineering/M2) to have a solid background in applied mathematics, Machine/Deep Learning, in particular generative models (diffusion models, flow matching), as well as in statistical signal/image processing and optimization. Knowledge in algorithms for sampling/variational inference and kernel methods will also be welcome. Interest or experience in differential geometry, optimal transport, and remote sensing applications will also be appreciated. The candidate is expected to have strong programming skills in Python (numpy/scipy/matplotlib...) and in Pytorch.

Environment The PhD thesis will take place at IMT Atlantique, Brest Campus, France, and is a 3-year (36 months) contract, expected to start around May 2026. The candidate will be part of the team of an AI Chair from the Brittany AI cluster SequoIA on generative modeling for inverse problems. The PI is part of the multidisciplinary research team (INRIA, IFREMER, IMT Atlantique, Univ. Brest) ODYSSEY which investigates the interplay between AI and inverse problems for ocean observation and reconstruction.

Contact To apply for this position, please send a detailed application including a cover letter, an up-to date CV, transcripts of grades and reference letters or any other element that you deem useful to:

- Lucas DRUMETZ, Associate Prof., IMT Atlantique at lucas.drumetz@imt-atlantique.fr. Google Scholar page
- Thierry CHONAVEL, Full Prof., IMT Atlantique at thierry.chonavel@imt-atlantique.fr.

Application deadline: January 31 2026

References

[1] J. M. Bioucas-Dias, Plaza *et al.*, "Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches," *IEEE JSTARS*, vol. 5, no. 2, pp. 354–379, 2012.

[2] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning.* MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[3] G. Daras, H. Chung, C.-H. Lai, Y. Mitsufuji, J. C. Ye, P. Milanfar, A. G. Dimakis, and M. Delbracio, "A survey on diffusion models for inverse problems," *arXiv preprint arXiv:2410.00083*, 2024.

[4] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," in *The Eleventh International Conference on Learning Representations*, 2023.

[5] M. Pourya, B. E. Rawas, and M. Unser, "Flower: A flow-matching solver for inverse problems," *arXiv preprint arXiv:2509.26287*, 2025.

[6] R. T. Chen and Y. Lipman, "Flow matching on general geometries," in *The Twelfth International Conference on Learning Representations*, 2024.

[7] C. Cheng, J. Li, J. Peng, and G. Liu, "Categorical flow matching on statistical manifolds," *Advances in Neural Information Processing Systems*, vol. 37, pp. 54 787–54 819, 2024.