



The Aramis Lab at Inria / Paris Brain Institute is recruiting a Master's student.

Generative Modeling & Optimal Transport for Medical Image Synthesis

Beginning of the internship: Spring 2026

Duration: 6 months

Financial support: regular internship gratification

Location: Institut du Cerveau, ARAMIS team, 187, rue Chevaleret, 75013 Paris

Keywords: Deep Learning, Optimal Transport, Sliced Wasserstein Distance, Medical Imaging, Image Synthesis, PyTorch

Collaboration: Joint supervision with Prof. Tom Vercauteren (King's College London)

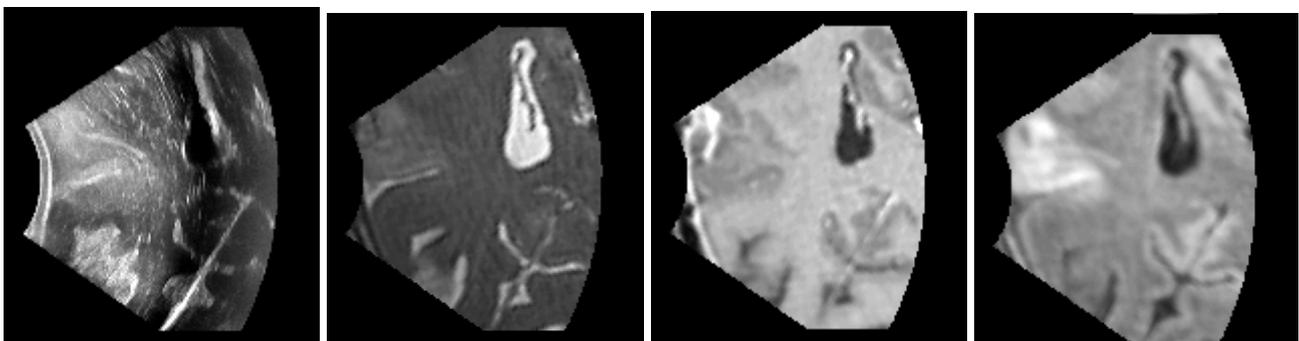
Contact

[Reuben Dorent](mailto:reuben.dorent@inria.fr): reuben.dorent@inria.fr

Context:

The **ARAMIS team** (Inria Paris) at the **Paris Brain Institute** is a computer science research group dedicated to developing computational models of brain diseases from multimodal patient data. The team is jointly affiliated with **CNRS, Inria, Inserm, and Sorbonne Université**, and operates within **Pitié-Salpêtrière Hospital**.

Within this environment, we are recruiting a highly motivated intern to work at the intersection of **deep learning, optimal transport, and medical image synthesis**, in close collaboration with **King's College London (KCL)**, under the supervision of **Reuben Dorent (Inria/ICM)** and **Tom Vercauteren (KCL)**.



From right to left: ultrasound, T2-MRI sequence, contrast-enhanced T1-MRI, FLAIR-MRI.
The goal is to synthesize one image modality from another using optimal transport and generative modeling.

Project Description:

Optimal Transport (OT) distances, and particularly **Wasserstein distances**, have become powerful tools for comparing probability distributions in imaging and generative modeling (e.g., [Villani, 2008](#); [Peyré & Cuturi, 2018](#)). However, computing full Wasserstein distances on high-dimensional medical images (such as 3D MRI and ultrasound) is computationally expensive.

To address this, **Sliced Wasserstein Distances (SWD)** approximate OT by projecting high-dimensional data onto 1D subspaces and aggregating 1D Wasserstein distances ([Rabin et al., 2011](#); [Nguyen et al., 2020](#), [Kolouri et al., 2019](#)). These approaches leverage the fact that the Wasserstein distance for one-dimensional probability densities has a closed-form solution computable in $O(n \log n)$. Yet, in the context of large 3D images where $n \approx 10^6$ voxels, minimizing these distances remains costly.

In this project, we aim to design **efficient SWD-based distances** for 3D medical image synthesis, with a focus on:

- Using **computationally efficient approximations of the 1D Wasserstein distance** instead of processing full voxel sets.
- Integrating these distances into **deep generative models** for realistic medical image translation ([Jiaqi He et al., 2024](#)).
- Evaluating how such approximations impact both **image quality** and **computational cost**, particularly for large 3D volumes typical of medical imaging.

You will be embedded in the **ARAMIS team** in Paris, working in **close collaboration with King's College London (Tom Vercauteren's group)**, with regular meetings and scientific exchanges.

What You Will Do:

As a member of the ARAMIS team, you will:

- Dive into 3D medical imaging data (e.g. MRI, ultrasound) and their specific computational challenges.
- Design and implement efficient sliced Wasserstein distance variants in 1D.
- Integrate these distances into deep learning pipelines for image synthesis (PyTorch).
- Experiment with and compare different OT-based losses and baseline generative models.
- Collaborate closely with researchers from Inria/ICM and King's College London.
- Contribute to scientific publications and open-source code.

Key Responsibilities:

- Develop and benchmark computer vision / deep learning methods for 2D/3D medical image synthesis.
- Study, implement, and possibly extend state-of-the-art OT and sliced Wasserstein methods from recent literature.
- Build and train models in PyTorch using SWD and related losses.
- Design and run quantitative and qualitative evaluations (metrics, visual assessment with domain experts).
- Targeted objective: an ICLR submission.

Required Skills & Knowledge:

- Motivated, curious, and eager to learn.
- Hands-on experience with Python and PyTorch (or an equivalent deep learning framework).
- Solid background in machine learning / deep learning and/or computer vision.
- Interest in medical imaging and real-world clinical applications.
- Comfortable reading and implementing methods from research papers.
- Optional: Basic knowledge or strong interest in Optimal Transport, Wasserstein distances, or related mathematical tools is a plus (but not strictly mandatory if you are ready to learn).