Probabilistic Non-Rigid Image Registration for Safe Brain Tumor Surgery

This PhD project is proposed within the PR[AI]RIE–PSAI doctoral programme, which promotes excellence in artificial intelligence and its interdisciplinary applications. In 2026, three PhD fellowships funded through Google philanthropy are offered within this programme.

If selected, the proposed project will be conducted within ARAMIS (Inria) at the Paris Brain Institute, under the supervision of Bertrand Mathon (HDR) and Reuben Dorent (PR[AI]RIE–PSAI junior fellow).

Doctoral project

Surgical resection is the critical first step for treating most brain tumors, and the extent of resection is the major modifiable determinant of patient outcome. To optimize outcome, surgeons must avoid causing neurologic deficits while leaving as little residual tumor as possible. To achieve these competing objectives, intraoperative ultrasound (iUS) has attracted attention as an affordable, real-time imaging technology that can be easily integrated into existing surgical workflows, especially compared with intraoperative Magnetic Resonance Imaging (MRI). However, iUS images are difficult to interpret due to low contrast between healthy and tumorous tissues and the presence of artifacts. For that reason, this intraoperative modality is not widely used by surgeons. In parallel, neuronavigation has helped considerably in providing intraoperative guidance to surgeons, allowing them to visualize the location of their surgical instruments relative to the tumor and critical brain structures visible in preoperative MRI. However, the utility of neuronavigation decreases as surgery progresses due to brain shift, which is caused by brain deformation and tissue resection during surgery, **leaving surgeons without guidance when and where it is most needed**, i.e., at the tumor margins. We believe brain shift compensation can be addressed by developing innovative probabilistic non-rigid registration algorithms that compute dense correspondences between preoperative MRI and intraoperative ultrasound images.

Given the extensive research in medical image registration, one might assume that existing multimodal approaches perform sufficiently well. Unfortunately, this is far from the truth. For example, in 2024 and 2025, we co-organized the international Learn2Reg[1] challenge at MICCAI, which included three registration tasks. Despite the participation of leading researchers and the use of diverse methods, none of the tasks were solved. On the one hand, traditional non-rigid registration methods rely on image-based metrics (e.g., mutual information) that are not well-suited to large multimodal domain gaps (e.g., iUS and MR) and are computationally expensive, making them impractical for intra-operative scenarios [1]. On the other hand, deep learning-based registration enables fast image registration (e.g., [2]) but often falls short in handling **large deformations** and maintaining **robustness** [3]. Alternatively, feature-based non-rigid registration allows for **real-time**, **interpretable** and **non-rigid** registration by focusing on a sparse set of local image feature correspondences. However, these methods tend to inadequately represent and match local content in **multimodal** images [3]. Moreover, creating dense displacement fields from sparse matches remains an open challenge. While nonlinear interpolation approaches using radial basis functions or splines are typically employed, they are not adapted to **changes in topology**. Finally, most of these techniques do not provide an **uncertainty quantification** that indicates where registration results can be relied upon. For all these reasons, simplistic, rigid registration methods, which provide robust results but are inherently not well-suited to large, complex deformations, are the most commonly used in practice.

The proposed doctoral project aims to develop *non-rigid*, *multimodal*, and *real-time* image registration methods tailored for *large deformations* with *topological changes* that explicitly predict *local uncertainty* to enable *informed decisions* by surgeons. To this end, we propose to develop new methods for **multimodal probabilistic feature matching** and **physically-inspired interpolation with uncertainty quantification** with four main objectives:

- **O1: Locate Driving Points at the Tumor Margin and Near Resection Boundary.** Discrete probabilistic displacement methods consider probabilistic distributions of deformations derived from driving point matches [4]. This project proposes to focus on driving points located at the key surgical boundaries visible in both iUS and MRI: the tumor, the dura mater (falx and tentorium), and the brain's surface (gyri and sulci). To extract such points, the project will leverage our previous research on deep-learning segmentation [5], domain adaptation [6], and image synthesis [7]. Specifically, the proposed project will 1) extend the previous brain structure and tumor segmentation framework (jSTABL) [5] trained with partially annotated publicly available datasets (BraTS with manual tumor annotations and ADNI with FreeSurfer annotations) to additionally segment the aforementioned brain structures; 2) perform cross-modality domain adaptation using the Mixture of Multimodal Hierarchical Variational Auto-Encoder (MMHVAE)

[1] `https://learn2reg.grand-challenge.org/`

synthesis framework [7] to transfer MRI-acquired knowledge to iUS; 3) validate the proposed approach on the ReMIND[2] data, a unique and open-source dataset containing preoperative MRI, intraoperative ultrasound and segmentations from 114 patients.

- **O2: Multimodal Feature Representations for Non-Rigid Registration.** This objective focuses on matching driving points obtained from O1 based on their local image contents. Instead of relying on handcrafted features (e.g., 3D SIFT-Rank [8], MIND [9]), which can lead to significant misalignment in MRI/iUS registration, this project proposes a self-supervised data-driven approach to learn deep representations adapted to image registration. The proposed approach will 1) learn informative and modality-specific deep representations of local content using self-supervised learning; 2) estimate the mutual information between pairs of multimodal representations using a discriminative approach proposed in [10]; 3) match driving points across multimodal images using the mutual information of the deep representations describing their local content as a similarity measurement. For 1), public, unpaired datasets (BraTS, ReSECT) will be used, as well as a unique dataset currently being acquired at Hôpital Pitié-Salpêtrière ($\sim$ 70 patients) containing ultrasound videos and preoperative MRI. For 2), the ReMIND dataset will be used, as it provides aligned pairs of iUS/MRI.
- **O3: Variational Inference for Uncertainty Registration Quantification.** Surgeons must know when and where image registration can be relied upon. This objective aims to improve surgeon interpretability by quantifying uncertainty derived from O2 driving-point matches, i.e., estimating posteriors. To address the inverse problem of estimating a dense displacement field from sparse matches, this project intends to make the first use of stochastic normalizing flows [11] to approximate posterior distributions from biomechanical model simulations, leveraging our previous work [12]. The proposed approach will start with simulating a training dataset of resection-induced deformations, including topological changes, using biomechanical finite element models. Subsequently, this project will explore approximating the posterior distribution using a variational inference framework that incorporates advanced conditional stochastic normalizing flows specifically designed to support topological changes. The use of latent diffusion models [13] could also be explored. To account for the high-dimensional parameter space typical of medical images, various architectures will be explored. Notably, the proposed framework provides Bayesian uncertainty quantification.
- **O4: Validate the proposed framework in clinically realistic neurosurgical scenarios.** The final objective of this project is to validate the proposed methods in close collaboration with clinicians and on large, clinically relevant datasets. The framework will be evaluated using the extended ReMIND database, which includes preoperative MRI, intraoperative ultrasound, and expert annotations of tumors, resection boundaries, and anatomical landmarks. Additionally, private data from Hôpital Pitié-Salpêtrière will be exploited. Quantitative evaluation will focus on registration accuracy at the tumor margin, robustness under large deformations, and the calibration and interpretability of uncertainty estimates. A panel of neurosurgeons from Pitié-Salpêtrière with different levels of experience will be included in the evaluation. This objective will ensure that the developed methods are not only technically innovative but also clinically meaningful and aligned with real intraoperative constraints.

A vibrant scientific, technological, clinical, and ethical environment

The selected candidate will work within the ARAMIS lab (`www.aramislab.fr`) and be located at the AI center of the Paris Brain Institute (ICM) (`https://institutducerveau-icm.org`). The institute, affiliated to Sorbonne Université, CNRS, Inserm, and AP-HP, is ideally located at the heart of the Pitié-Salpêtrière hospital. The ARAMIS lab, which is also part of Inria, is devoted to the design of computational, mathematical, and statistical approaches for the analysis of multimodal patient data in brain disorders, with an emphasis on imaging data.

A strength of this project is its multidisciplinary aspect, combining clinical neuroscience, imaging, and artificial intelligence. The PhD will be co-supervised by Bertrand Mathon, neurosurgeon at Hôpital Pitié-Salpêtrière and PI of the BRIGHT team at ICM, and Reuben Dorent, permanent Inria researcher in ARAMIS and junior fellow of PR[AI]RIE-PSAI. Bertrand Mathon obtained his HDR in 2024 and has received awards from the French National Academy of Medicine (2017), the French National Academy of Surgery (2021), and AP-HP (2023). Reuben Dorent completed his PhD in 2022 at King's College London, followed by a postdoc at Harvard University and a Marie Skłodowska-Curie fellowship at Inria Saclay.

Beyond the ICM/Inria research environment, the proposed project would greatly benefit from the rich multidisciplinary PR[AI]RIE-PSAI ecosystem as it is particularly well aligned with its dual ambition. From a methodological perspective, the project tackles key AI challenges in multimodal learning, representation learning, and probabilistic modeling, with the aim of developing tools for computer-aided decision support.

[2] `https://www.cancerimagingarchive.net/collection/remind/`

From an application perspective, it targets a relevant clinical challenge by aiming to improve image-guided brain surgery, reduce the risk of postoperative neurological deficits, and ultimately improve patient outcomes.

At the same time, PR[AI]RIE-PSAI would provide an excellent training environment for the PhD candidate. Through its seminars, summer schools, and interactions with internationally recognized researchers, the student would benefit from broad exposure to state-of-the-art AI methods and their diverse applications. The ecosystem would also offer valuable opportunities to strengthen scientific culture, develop a strong interdisciplinary network, and gain exposure to innovation through possible collaborations with industry and support for entrepreneurship. As such, the project is not only relevant to PRAIRIE's scientific objectives, but would also provide an excellent framework for training a future researcher at the interface of AI and healthcare.

Non-discrimination, openness, and transparency / Non-discrimination, ouverture et transparence

All partners of PR[AI]RIE-PSAI commit to supporting and promoting equality, diversity, and inclusion within their communities. We encourage applications from individuals with diverse backgrounds, which we will ensure are considered through an open and transparent recruitment process. *L'ensemble des partenaires de PR[AI]RIE-PSAI s'engagent à soutenir et promouvoir l'égalité, la diversité et l'inclusion au sein de ses communautés. Nous encourageons les candidatures issues de profils variés, que nous veillerons à sélectionner via un processus de recrutement ouvert et transparent.*

Your profile

- Master or engineering degree with a specialization in machine learning
- Strong interest in medical applications
- Good programming skills in Python
- Knowledge in digital image processing and medical imaging
- Good writing skills
- Good relational and communication skills

Application process

Applications for this PhD project should be submitted to the supervisors at reuben.dorent@inria.fr. **The application deadline is May 10**. Candidates are required to submit a complete application including: (i) a curriculum vitae, (ii) a one-page motivation letter describing their interest in the proposed research topic, their ambitions, and the relevance of their background to the project, and (iii) copies of their most recent academic diplomas. The selection results will be communicated between May 30 and mid-June at the latest.

References

[1] Klein, A., ..., Parsey, R. V, "Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration", in *Neuroimage* **46(3)**, 786-802 (2009).

[2] Mok, T. C., Chung, A., "Fast symmetric diffeomorphic image registration with convolutional neural networks", in *CVPR* (2020)

[3] Hering, A., ..., Heinrich, M. P., "Learn2Reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning", in *IEEE Trans. Med. Imaging*, **42(3)**, 697-712 (2022).

[4] Heinrich, M. P., "Closing the gap between deep and conventional image registration using probabilistic dense displacement networks", in *MICCAI* (2019).

[5] Dorent, R., ..., Vercauteren, T. "Learning joint segmentation of tissues and brain lesions from task-specific heteromodal domain-shifted datasets", in *Med. Image Anal.* **67**, 101862 (2021).

[6] Dorent, R., ..., Vercauteren, T., "CrossMoDA 2021 challenge: Benchmark of cross-modality domain adaptation techniques for vestibular schwannoma and cochlea segmentation", in *Med. Image Anal.* **83**, 102628 (2023).

[7] Dorent, R., .., Wells, W., "Unified Cross-Modal Medical Image Synthesis with Hierarchical Mixture of Product-of-Experts", in *IEEE Trans. Pattern Anal. Mach. Intell.* **48** (2026)

[8] Toews, M., ..., Arbel, T., "Feature-based morphometry: Discovering group-related anatomical patterns", in *Neuroimage* **49(3)**, 2318-2327 (2010).

[9] Heinrich, M. P., ..., Schnabel, J. A., "MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration", in *Med. Image Anal.* **16(7)**, 1423-1435 (2012).

[10] Dorent, R., ..., Wells, W., "Connecting Jensen–Shannon and Kullback–Leibler Divergences: A New Bound for Representation Learning", in *NeurIPS*, (2025).

[11] Hagemann, P., ..., Steidl, G., "Stochastic normalizing flows for inverse problems: A Markov chains viewpoint", in *SIAM/ASA J. Uncertainty Quantif.*, **10(3)**, 1162-1190 (2022).

[12] Assis, T., ..., Dorent, R., "Deep Biomechanically-Guided Interpolation for Keypoint-Based Brain Shift Registration", in *COLAS* (2025).

[13] Rout, L., ..., Chu, W. S., "Beyond first-order tweedie: Solving inverse problems using latent diffusion", in *CVPR* (2024).