Application to be sent to Christa Rachelle Joseph and Scarlett Weale:
christa-rachelle.joseph@utoulouse.fr scarlett.weale@utoulouse.fr
Deadline: February 27st, 2026

EUR CARe PhD program pre-proposal

(2 pages maximum)

PhD Director 1 : *Emmanuelle Claeys ; emmanuelle.claeys@irit.fr*
PhD Director affiliation: Institut de Recherche en Informatique de Toulouse (IRIT), *MISFIT (Machine Learning Integrity & Safety Fairness Impact Trust)*
PhD co-Director 2 : *Rodrigo Moreno; rodmore@kth.se*
PhD co-Director affiliation: *Professor, KTH Royal Institute of Technology* (Stockholm,Sweden) - Division of Biomedical Imaging
PhD co-Director 3: *Ken Soléakhéna ; ken.soleakhena@iuct-oncopole.fr*
PhD co-Director affiliation: Cancer Research Center of Toulouse (CRCT) / Institut Universitaire du Cancer de Toulouse – Oncopole Claudius Regaud

Research project title: *Preference-Constrained Multimodal RLHF to Trigger Offline Adaptive Radiotherapy Under Clinical Resource Constraints*

Research program abstract (max 500 words):

Offline adaptive radiotherapy (offline ART) is an adaptive radiotherapy approach in which the treatment plan is adjusted between treatment sessions, rather than in real time during irradiation. A key aspect of offline ART is that the patient's anatomy evolves over days or weeks (e.g., weight loss, tumor shrinkage, organ displacement or repletion…). It consists of observing these anatomical changes and recomputing a new treatment plan before subsequent sessions. Clinically, patient management begins with an initial plan based on a reference computed tomography (CT) scan. During treatment, control imaging (such as CBCT, repositioning CT, or MRI [5]) is performed. An offline analysis compares the current anatomy with the initial plan, and replanning is performed if necessary, after which subsequent fractions are adapted accordingly. A crucial question in this workflow is: "**Does the observed anatomical difference justify replanning?**" **This decision point is clinically of high-impact because it directly trades off avoiding organs at risk (OAR) overdoses or loss of dose coverage on the tumor target against feasibility (time, staffing, and machine availability).**

Today, this decision relies on heuristic thresholds (e.g., volume variation, dose deviation, distance to OAR) and on clinical judgment from radiation oncologists. Offline ART can be modelled as a reinforcement learning (RL) problem. At fraction t, an agent observes anatomical differences, dosimetric indicators, and a history of past decisions. The agent must then select an action at each fraction, either discrete (do nothing, minor replanning, full replanning) or continuous (adjust margins, modify constraints, trigger a new CT acquisition). The proposed RL model must integrate multimodal information, including medical images, textual descriptions, and tabular data. **The novelty is to learn a *human-aligned* triggering policy from real clinical preferences rather than encoding a fixed rule-set, enabling consistent decisions across teams and sites.**

A second challenge will be to train a reward model that approximates "**what an expert considers a good adaptation decision**." This corresponds to reinforcement learning with human feedback (RLHF). Optimization through RLHF aims to avoid unnecessary replanning while not missing critical adaptations, while keeping the clinician in the loop. **Clinically, this is expected to reduce preventable toxicities and missed adaptations while also decreasing avoidable replans**.

Foundation models are transformer-based models and are pre-trained at a large scale on heterogeneous data and adapted to specific tasks through fine-tuning or contextual conditioning, providing a shared representational backbone reusable across radiology applications [2,3,4]. **A further novelty is to evaluate whether these reusable multimodal backbones can improve data efficiency and generalization of offline ART triggering across centers.**

Hierarchical Sequential Transduction Units (HSTU) have also shown promising state-of-the-art results for sequential recommendation tasks, as well as scalability to extremely large datasets [1]. However, there is still limited feedback on

their application to medical reinforcement learning. **This thesis will provide one of the first systematic comparisons of sequential recommender architectures (including HSTU) versus foundation-model-based approaches for a safety-critical, longitudinal radiotherapy decision task.** Importantly, the project is **already supported by access to relevant retrospective clinical data** within the current clinical/research environment, enabling immediate feasibility and early prototyping.

Describe in 50 words max for each how this project fits the 3 defining criteria of the CARe graduate programme:

Relation to CARe topics of Cancer, Ageing and/or Rejuvenation (≤50 words)

This project aligns with CARe by developing **human-centered, reliable and clinically relevant AI**. By integrating medical expertise directly into deep RL through **RLHF**, it improves trust, robustness and clinical acceptability of decision support for adaptive radiotherapy, particularly relevant for aging populations, with cumulative diseases burden and history of cancer irradiations.

Multidisciplinary aspect (≤50 words)

Interdisciplinary by design: **AI/RLHF** (E Claeys, IRIT), **medical imaging & AI** (R Moreno, KTH), **medical physics & access to retrospective data** (S Ken, IUCT-OCR), and **radiotherapy expertise** with C Massabeau, MD (IUCT-OCR) specialized in 4D-CT lung cancer. The project aims at bridging AI, machine learning, imaging, and clinical decision-making.

International and/or industrial aspect(s) (≤50 words)

Strong international dimension through co-supervision with **KTH (Stockholm, Sweden)** and planned cross-center generalization studies. The project builds on **existing IRIT/IUCT-OCR (Toulouse, France) collaborations** in multidisciplinary radiotherapy/imaging research (e.g., IABM 2026 two accepted abstracts) and aims to extend this established collaboration network to a broader European scale. In addition, future discussions with collaborators at Karolinska Hospital could enable access to complementary local datasets, if of interest, to further test cross-center generalization.

5 keywords in line with EUR CARe:

Adaptive radiotherapy; Offline replanning trigger; Multimodal learning; Reinforcement learning from human feedback (RLHF); Workflow-constrained decision support

5 references of the teams, highlighting the co-signatory students:
(Last 5 years)

[1] Reinforcement Learning with Human-Guided Reward Model for Improved MRI/4D-CT Deformable Image Registration, U Harlouchet, Z Kheil, **E Claeys**, **S Ken** *ESTRO 2026, Abstract 5066*

[2] Dynamic Allocation Optimization in A/B-Tests Using Classification-Based Preprocessing, **E Claeys**, P Gancarski, M Maumy-Bertrand, H Wassner *IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 1, pp. 335-349, 1 Jan. 2023, doi: 10.1109/TKDE.2021.3076025*

[3] Photon-counting detector CT and energy-integrating detector CT for trabecular bone microstructure analysis of cubic specimens from human radius, B Klintström, L Henriksson, **R Moreno**, A Malusek, Ö Smedby, …, *Eur Radiol Exp. 2022 Jul 27;6(1):31. doi: 10.1186/s41747-022-00286*

[4] A deep learning model for brain age prediction using minimally preprocessed T1w images as input, C Dartora, A Marseglia, G Mårtensson, G Rukh, J Dang, JS Muehlboeck, LO Wahlund, **R Moreno**, …,*Front Aging Neurosci. 2024 Jan 8:15:1303036. doi: 10.3389/fnagi.2023.1303036*

[5] Clinical outcome and deep learning imaging characteristics of patients treated by radio-chemotherapy for a "molecular" glioblastoma. C Zerbib, L Robinet , **S Ken**,..., E Cohen-Jonathan Moyal . *Oncologist. 2025 Jun 4;30(6):oyaf127. doi: 10.1093/oncolo/oyaf127*.

Appendix

Additional References

[1] Zhai, et al., “Actions Speak Louder than Words: Trillion-Parameter Sequential Transducers for Generative Recommendations”, https://arxiv.org/abs/2402.17152 ,ICML 2024

[2] Attention Is All You Need, Ashish Vaswani and Noam Shazeer and Niki Parmar and Jakob Uszkoreit and Llion Jones and Aidan N. Gomez and Lukasz Kaiser and Illia Polosukhin, 2023

[3] Model-based Reinforcement Learning: A Survey Thomas M. Moerland, Joost Broekens, Aske Plaat, Catholijn M. JonkerThomas M. Moerland, Joost Broekens, Aske Plaat, Catholijn M. Jonker, 2022

[4] Y. He et al., "Foundation Model for Advancing Healthcare: Challenges, Opportunities and Future Directions," in IEEE Reviews in Biomedical Engineering, vol. 18, pp. 172-191, 2025, doi: 10.1109/RBME.2024.3496744.

PhD Directors / Institutions Overview

Emmanuelle Claeys : Emmanuelle Claeys is an Associate Professor at Université Paul Sabatier and a member of the IRIT laboratory. She leads the newly established MISFIT team (Machine Learning Integrity & Safety Fairness Impact Trust). She also heads the AOC project team, a collaboration in mathematics and computer science dedicated to artificial learning. Her primary research area is reinforcement learning. She collaborates with LAAS and IMT on biomechanical learning problems for robotics. She is also a member of an ANITI chair focused on fair and robust machine learning. In addition, she supervises several PhD theses on data generation techniques and reinforcement learning.

The Institut de Recherche en Informatique de Toulouse (IRIT) is one of the largest Joint Research Units at the national level and a major pillar of research in the Occitanie region. With approximately 600 members, including permanent and non-permanent staff, and around 100 external collaborators, IRIT plays a structuring role in computer science research and its applications in digital technologies at both regional and national levels. Due to its multi-institutional supervision (CNRS and Toulouse universities), strong scientific impact, and interactions with other disciplines, IRIT is a key actor in the French computer science research landscape. Within IRIT, the MISFIT team focuses on the major challenges of trustworthy artificial intelligence, aiming to strengthen safety, fairness, reliability, and trust in machine learning systems.

Rodrigo Moreno : Rodrigo Moreno is a Professor and researcher specializing in computational medical imaging at the KTH Royal Institute of Technology. The core of his research focuses on designing mathematically grounded computational tools for the analysis of medical images acquired from multiple imaging modalities. These tools also aim to improve image acquisition protocols. The efficiency of the proposed methods is a central concern, as their ultimate goal is to support clinicians in everyday medical practice. Rodrigo Moreno places strong emphasis on the translation of fundamental research advances into concrete clinical applications, which involves close collaboration with industrial partners. His work notably includes the use of advanced artificial intelligence methods to address complex image analysis problems, as well as the optimization of acquisition and processing pipelines for advanced MRI modalities.

The KTH Royal Institute of Technology is a leading university located in Stockholm, Sweden. KTH is internationally recognized for its excellence in engineering sciences, applied mathematics, computer science, and medical technologies, as well as for its strong collaborations with industry and university hospitals.

Ken Soléakhéna : Soléakhéna Ken, PhD, HDR, has been working for over 15 years within the Medical Physics Department of IUCT-Oncopole. Within the Cancer Research Center, his main responsibilities involve defining the most relevant imaging modalities to be acquired for various clinical trials, which requires close collaboration with radiologists, oncologists, radiation oncologists, and medical imaging technologists. He has supervised the design of several national, multi-institutional clinical trial protocols, ranging from neuro-oncology studies to a whole-body imaging trial for chronic lymphocytic leukemia (CLL). The latter relied exclusively on optimized MRI acquisition, a non-conventional imaging modality used for CLL monitoring.

Oncopole Claudius Regaud, a comprehensive center for cancer care, research, and education, brings together the expertise of more than 2,300 professionals on a single site in Toulouse and is accredited as a Comprehensive Cancer Center. It combines state-of-the-art clinical facilities for cancer treatment with world-class research infrastructure within an integrated campus that includes both public and private stakeholders, as well as industrial partners. Oncopole will provide the project with a multimodal imaging acquisition protocol specifically designed for the longitudinal follow-up of patients with chronic lymphocytic leukemia. The available dataset will include, for each patient, whole-body MRI, incorporating specific morphological observations of patient anatomy, along with diffusion-weighted sequences enabling the measurement of free-water diffusion markers between cells.