

Job Title: Post-Doctoral Researcher in Efficient Foundation Models.

Location: ADAPT Research Centre, Trinity College Dublin, Ireland

Duration: 18 Month Specific Purpose Fixed Term

Application Deadline: 10th March 2025

Start Date: 10th June 2025 (or sooner if available)

Salary scale: €44,847 - €69,692

We are seeking a talented post-doctoral researcher to join our research team to innovate in the area of **Efficient Foundation Models (EFMs)**. This project aims to develop novel techniques that enable adaptation of FMs to narrow domains and reduce inference costs under resource constraints. The research will focus on three key areas:

- **KT 1: Neural Composition:** Investigating modular and compositional architectures to enable flexible domain adaptation by combining the strengths of different models.
- **KT 2: Adaptive Computation:** Exploring techniques such as dynamic pruning / routing, and Mixture-of-Experts models to maximize inference efficiency, without compromising accuracy.
- **Reinforcement Learning and Optimization:** Apply principled machine learning methods to achieve breakthroughs in KT 1 and KT 2.

This position offers the opportunity to work at the forefront of AI and NLP research, contributing to high-impact publications, and collaborating with leading academic and industry partners.

Responsibilities

- Conduct cutting-edge research on efficient domain adaptation and efficient inference of foundation models (principally language models).
- Develop and systematically evaluate novel approaches for neural composition and adaptive computation, drawing on techniques from multi-objective reinforcement learning and optimization (e.g. evolutionary algorithms).
- Collaborate with interdisciplinary teams, including machine learning researchers and domain experts.
- Publish research findings in top-tier conferences and journals.
- Contribute to the development of open-source tools and frameworks for efficient model training and inference.
- Lead on project documentation and reporting activities

Qualifications

Essential:

- PhD in Computer Science, Machine Learning, Natural Language Processing, or a related field.
- Strong publication record in relevant AI/ML/NLP venues (e.g., NeurIPS, ICML, ACL, EMNLP).
- Proficiency in deep learning frameworks (e.g., PyTorch, TensorFlow).
- Experience training and deploying language models (e.g., transformer-based architectures such as BERT, GPT, T5).
- Strong understanding of machine learning fundamentals; such as MORL/EA, NLP, transfer learning methods, and design of experiments.
- Strong programming skills in Python.
- Excellent problem-solving and analytical skills.

Desirable:

- Familiarity with techniques for efficient training of LLMs; such as quantization, pruning, and distillation.
- Familiarity with techniques for efficient inference, such as Mixture-of-Experts models, adaptive computation, etc.
- Understanding of recent trends in the areas of multi-agent systems and reasoning with LLMs.
- Prior experience with large-scale computing environments (e.g., distributed training, HPC clusters).
- Experience collaborating with industry or applied research projects.

Application Process

Interested candidates should submit the following documents to Stephen Rigney at stephen.rigney@adaptcentre.ie

- A cover letter detailing their research experience and how it aligns with the project.
- A full academic CV, including a list of publications.
- Contact details of at least two academic referees.

We encourage applications from diverse backgrounds and are committed to creating an inclusive research environment.

For general inquiries about the position, please contact Patricia Buffini at patricia.buffini@adaptcentre.ie

We look forward to your application!