Research Intern - LLM Fine-tuning and Agentic AI

Technology Innovation Institute (TII) | AI and Digital Science Research Center | Abu Dhabi, UAE

Type: Internship, full-time, on-site | Duration: 6 months

The Institute

Technology Innovation Institute (TII) is a publicly funded research institute based in Abu Dhabi, United Arab Emirates. It brings together scientists, engineers, mathematicians and researchers from around the world to turn hard problems into pioneering research and technology prototypes.

The Team

The AI and Digital Science Research Center is the home of the Falcon family of open large language models. The team covers the full model development stack, from pretraining data to post-training, evaluation and release, and is now pushing further into reasoning and agentic capabilities.

Role Summary

We are looking for motivated interns to join the Falcon post-training effort. You will work hands-on with large language models, contributing to fine-tuning, reinforcement learning and agentic systems, under the guidance of senior researchers. The internship is research-oriented, with the goal of producing results that feed into Falcon models and, where relevant, publications or open-source releases.

What You Will Work On

- Supervised fine-tuning and preference optimization (e.g., SFT, DPO) of LLMs for reasoning, coding and instruction following.
- Reinforcement learning post-training of LLMs (e.g., GRPO-style methods) with verifiable rewards.
- Agentic AI: training and evaluating models for multi-step reasoning, planning and tool use in interactive environments (e.g., code and terminal sandboxes).
- Building data pipelines for training and synthetic data generation.
- Designing evaluation pipelines and running benchmarks for fine-tuned and agentic models.
- Reproducing recent papers, running ablations and reporting findings clearly to the team.

Requirements

- Currently enrolled in an MSc or PhD in Computer Science, AI/ML, Mathematics or a related field (strong final-year BSc students may also apply).
- Solid understanding of deep learning, transformers and how LLMs are trained.
- Strong programming skills in Python and hands-on experience with PyTorch.
- Familiarity with the Hugging Face ecosystem (Transformers, Datasets, TRL or similar).
- Comfortable working on Linux, with git and GPU environments.
- Good communication skills in English and ability to work in a team.

Nice to Have

- Experience with RL for LLMs (PPO, GRPO, DPO and variants).
- Experience with distributed training (DeepSpeed, FSDP, Megatron) or inference engines (vLLM, SGLang).
- Experience building or training LLM agents, tool-use pipelines or sandboxed environments.

Document Classified as: Internal

- Publications at ML/NLP venues (e.g., NeurIPS, ICML, ICLR, ACL, EMNLP) or meaningful open-source contributions.

What We Offer

- Work on the Falcon models with access to large-scale compute.
- Close mentorship from senior researchers on the Falcon team.
- Opportunity to contribute to publications and open-source model releases.

How to Apply

Send your CV, a short statement of interest (topics you want to work on and your availability), and links to your GitHub, Google Scholar or relevant projects to:

Reda Alami, Senior Research Scientist

reda.alami@tii.ae

Please use the subject line: "Internship Application - LLM Fine-tuning and Agentic AI".

Document Classified as: Internal