

Research Engineer or Postdoctoral Position

Post-Training Generative Models for Education at Scale

Flowers AI & CogSci Lab, Inria Bordeaux (France)

Supervision: [Pierre-Yves Oudeyer](#)

Duration: up to 3 years (flexible, start as soon as possible)

Working language: English (French useful but not required)

In short

This position offers a rare combination:

- **Frontier post-training research** (SFT, preference optimization, RLHF/RLVR, knowledge tracing, agentic LLMs) applied to a real educational problem.
- **Real-world scale:** thanks to our industrial collaborator [EvidenceB](#), we have access to large-scale educational software (and associated data) AdaptivMaths and MIA Seconde, and others, which are deployed across French primary, middle, and high schools and reach tens of thousands of students.
- **World-unique data:** we have access to very large educational datasets with rich semantic and multimodal information about authentic student–content interaction; this is among the most extensive of their kind for training education-specific generative models.
- **Large-scale compute:** the position has access to a large H100 allocation on the Jean Zay supercomputer, enabling training at scale rather than at toy scale.
- **Publication and open science:** targets include NeurIPS, ICML, ICLR, AIED, LAK, EDM. Datasets, models, and software will be released as open-source digital commons.

If you want to train models (not just prompt them) at a scale and on data that almost no academic position offers, this is a strong match.

Project: GAIMHE

GAIMHE (Generative AI for Hybrid Mathematics Education) is a large-scale research and innovation project funded by Bpifrance, in partnership with EvidenceB, Café pédagogique, and ClassCode. It addresses a foundational tension in educational AI: intelligent tutoring systems deliver pedagogically grounded personalization but require heavy manual content authoring, while modern LLMs are generative and flexible but pedagogically ungrounded, weak at sustained long-term personalization, and computationally expensive. GAIMHE develops hybrid architectures that combine both — and the deployment partnership means innovations can be validated in real classrooms at unusual scale. In alignment with

open-science principles, and through partnership with Région Île-de-France, major project outputs will be released under open-source licenses.

Mission

The mission is intentionally flexible and shaped around the candidate's strengths and interests. Concrete directions we are excited about include:

- **Post-training for pedagogical alignment:** how to fine-tune compact LLMs to generate exercises, hints, and feedback that satisfy structured pedagogical constraints (SFT, DPO/GRPO, RL with verifiable rewards).
- **Compact student models:** training small, efficient models that predict learning outcomes and learning progress across pedagogical graphs, extending recent work such as MAGELLAN [3].
- **Agentic curriculum design:** LLM agents that autonomously generate, evaluate, and curate educational content while remaining aligned with cognitive-science-based learning principles.
- **Education-specialized foundation models:** leveraging the partner platforms' large-scale, multimodal, longitudinal interaction data to specialize generative models for the domain.

PhD-holders pursuing a publication-led postdoc trajectory and engineering-minded MSc graduates pursuing an implementation-led research-engineer trajectory are both welcome; the precise mission and priorities will be tuned with the candidate.

Who we're looking for

This position is for someone who **builds and trains** models, not just uses them.

Essential

- MSc or PhD in ML/AI. Masters MVA/IASD/Mind/etc, or equivalent international programs especially welcome.
- Demonstrated hands-on experience implementing and debugging modern post-training pipelines (SFT, DPO/PPO/GRPO, RLHF, RLVR). We will ask you to describe a training run you personally led.
- Fluency in PyTorch and the HuggingFace ecosystem at the level of modifying training code, not only configuring it.
- Comfort reading and reproducing recent papers on post-training, RL fine-tuning, and LLM agents.
- Strong written and spoken English. French is useful for interaction with educational partners but is not a requirement.
- Ability to work collaboratively in an interdisciplinary research team.

Not a fit

Candidates whose LLM experience consists primarily of prompt engineering, RAG pipelines, or agent orchestration over closed APIs without underlying training experience. This role is centered on model training; please do not apply if that is not where your strengths lie.

Valued

- Experience with multi-GPU and distributed training, Slurm-based HPC clusters (Jean Zay or equivalent).
- Background in knowledge tracing, learning analytics, cognitive science, or the science of learning.
- Open-source contributions, particularly to training or RL libraries.
- Publications at NeurIPS, ICML, ICLR, AIED, LAK, EDM, or equivalent venues.
- Experience deploying ML systems in production.

What we offer

- A **large H100 allocation on Jean Zay** — the ability to actually train, not just prototype.
- Access to **world-unique large-scale educational datasets** with rich semantic and multimodal information about authentic student–content interactions.
- A **25-person interdisciplinary lab** at the crossroads of AI and cognitive science, pioneering curiosity-driven learning, agentic LLMs, and educational AI — with daily collaboration across these threads (see [research program](#)).
- **Flexible scope**: research-engineer profiles can lead substantial engineering with co-authorship on resulting publications and the option to convert to a PhD; postdoc profiles can pursue independent publication-led research at top ML and EdTech venues.
- A strong **open-science culture**: code, models, and datasets released under open-source licenses.

Team context

The Flowers AI & CogSci lab at Inria Bordeaux does research at the crossroads of AI and cognitive sciences, focusing on studying curiosity-driven open-ended learning in AI and humans (see [research program here](#)). Our interdisciplinary team (25 members) has been pioneering the design of curiosity-driven learning algorithms in [AI](#) and [robotics](#), recently pioneering the design of [agentic generative AI systems](#) that learn to solve [self-generated goals](#) in interactive environments, ranging from video games to [ARC-AGI](#) to [self-improvement in coding skills](#). We have also developed new theories of human curiosity, like the [Learning Progress Hypothesis](#). Our main domains of application are [educational technologies](#) (related to this position) and [assisted scientific discovery](#). This position will involve daily interactions with all members of the team working on these dimensions.

How to apply

Send to pierre-yves.oudeyer@inria.fr with **[application]** in the subject line:

1. **CV**
2. **Motivation letter**: what draws you to this project and which scientific directions interest you most.
3. **A short technical write-up (≈1 page)** describing the most complex model training run you have personally led: what was trained, dataset and preprocessing, architecture and loss choices, optimizer and hyperparameter decisions, what broke and how you debugged it, and how you evaluated the result. *This is a required part of the application.*
4. **Links to code repositories**: especially training code you have written.
5. **Reports or write-ups of previous research projects** (not necessarily on this topic).
6. **Recent academic transcripts**, if relevant.

Applications are reviewed on a rolling basis.

References

[1] Shen, S., Liu, Q., Huang, Z., Zheng, Y., Yin, M., Wang, M., & Chen, E. (2024). A Survey of Knowledge Tracing. *arXiv preprint arXiv:2305.12345*. (NB: please verify the arXiv ID before publication — looks like a placeholder in the original.)

[2] Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., & Sohl-Dickstein, J. (2015). Deep Knowledge Tracing. *NeurIPS*, 505–513.

[3] Gaven, L., Carta, T., Romac, C., Colas, C., Lamprier, S., Sigaud, O., & Oudeyer, P.-Y. (2025). MAGELLAN: Metacognitive predictions of learning progress guide autotelic LLM agents in large goal spaces. *ICML 2025*. arXiv:2502.07709.

[4] Clément, B., Sauzéron, H., Roy, D., & Oudeyer, P.-Y. (2024). Improved performances and motivation in intelligent tutoring systems: combining machine learning and learner choice. *arXiv preprint arXiv:2402.01669*.