CRITEO
AI Lab

PhD proposal: self-improving AI agents for recommendation

Contact : Alain Rakotomamonjy, a.rakotomamonjy@criteo.com, Alberto Lumbreras, a.lumbreras@criteo.com, Patrick Gallinari, p.gallinari@criteo.com
Location: Criteo AI Lab, Paris
Candidate profile: Master degree in computer science or applied mathematics, Engineering school. Background and experience in machine learning. Good technical skills in programming.
How to apply: please send a cv, motivation letter, grades obtained in master, recommendation letters when possible to the contacts
Start date: January 2026 for three years
Keywords: deep learning, language models, generative models, information retrieval, recommendation

Context and Research Motivation
As part of its ongoing transformation into an agentic-ready platform, Criteo is spearheading the integration of agentic AI across its full portfolio. These systems are already being deployed to automate internal operations, assist clients in the management and optimization of advertising campaigns, and to power personal shopping agents—autonomous assistants that act on behalf of end-users. These agents must reason, remember, and act autonomously in environments characterized by uncertainty, variability, and scale.
To fulfill this vision, one of the most pressing challenges is adaptability. Our agents must function across an extremely heterogeneous client base — each with unique product catalogs, optimization targets, and interface constraints while interacting with users and inferring their intents.

Scientific Objectives
The objective of the PhD is to explore adaptation strategies to multiple and heterogeneous environments and user segments for an agentic system. In our setting these environments might correspond to different partners characterized by their own catalog, objective and strategy while user segments refer to user preferences or needs. We will restrict our scope to language-only agents and emphasize practical assistant scenarios.

In most scenarios, adaptation to new environments and to user intents shall leverage simple and computationally costless strategies, while being able to adapt for scarce data contexts available for these new settings. Adaptation places a significant demand on the system's memory, which must be more than a static repository of facts. It must be an adaptive memory system, capable of restructuring and reprioritizing information as the user's context evolves. Therefore, self-adaptation is intrinsically linked to memory management. The goal is to endow the agent with the ability to learn how to manage its own memory in response to a changing environment and user. The PhD will start to investigate different memory strategies and their potential for handling adaptation to new environments and to user interaction. We will explore mechanisms for the agent to develop learned policies for memory operations. Key research questions include:

- Learned Retention and Forgetting: How can an agent learn what information is critical to retain versus what is obsolete and should be forgotten or archived?
- Adaptive Retrieval Strategies: Can an agent learn the most effective way to query its memory? We will explore how the system can dynamically choose between different retrieval methods (e.g., vector-based RAG, evolving LLM context), based on the task.
- Automated Memory Summarization: How can the system "reflect" on its interaction history to create higher-level insights? We will investigate techniques for the agent to periodically

summarize streams of memories into more abstract knowledge (e.g., consolidating multiple shopping interactions into a persistent preference like "user prefers sustainable brands").

Adaptation mechanism shall also be an element contributing to the planning mechanism of the agent**:** how can an agent make decisions when the goal is weakly defined, the feedback is sparse, and the environment varies by client? This is particularly relevant in domains like travel planning or multi-product recommendations, where a "one-size-fits-all" approach is neither feasible nor desirable. To complement memory-based methods, off-line reinforcement learning strategies could be considered.

Methodology and Roadmap
The project will begin with a literature review and benchmarking of existing memory and adaptation frameworks. Building on this foundation, we will implement proof-of-concept agents in both synthetic and real-world commerce use cases. Over time, the student will work on designing a unified, modular framework that integrates memory, adaptation, and planning for agentic commerce, and benchmark it on internal metrics such as engagement, regret, and generalization gap.

Research Environment

The PhD will be hosted by CRITEO AI-Lab in Paris. The student will collaborate closely with experienced research scientists and engineers, and have access to internal datasets, production-grade infrastructure, and experimentation platforms. There will also be opportunities to contribute to real product initiatives and to publish in top-tier conferences.

Candidate Profile
We are looking for a motivated researcher with a strong foundation in machine learning, natural language processing, applied maths. Familiarity with large language models, transformers, reinforcement learning, or continual learning will be considered a strong asset. Above all, we are seeking someone who is excited by the challenge of bringing intelligent agents to life in practical, high-impact applications.

References

- Chhikara, P., Khant, D., Aryan, S., Singh, T., & Yadav, D. (2025). Mem0: Building Production-Ready AI Agents with Scalable Long-Term Memory. arXiv preprint arXiv:2504.19413.
- Goldie, A., Mirhoseini, A., Zhou, H., Cai, I., & Manning, C. D. (2025). Synthetic Data Generation & Multi-Step RL for Reasoning & Tool Use (SWiRL). arXiv preprint arXiv:2504.04736.
- Packer, C., Wooders, S., Lin, K., Fang, V., Patil, S. G., Stoica, I., & Gonzalez, J. E. (2023). MemGPT: Towards LLMs as Operating Systems. arXiv preprint arXiv:2310.08560. arXiv+1
- Salama, Rana; Cai, Jason; Yuan, Michelle; Currey, Anna; Sunkara, Monica; Zhang, Yi; Benajiba, Yassine (2025). MemInsight: Autonomous Memory Augmentation for LLM Agents. arXiv preprint arXiv:2503.21760
- Xu, Wujiang; Liang, Zujie; Mei, Kai; Gao, Hang; Tan, Juntao; Zhang, Yongfeng (2025). A-MEM: Agentic Memory for LLM Agents. arXiv preprint arXiv:2502.12110.