MANDANA SAMIEI

Montréal, Canada | mandanasmi.github.io | samieima@mila.quebec | GitHub | Google Scholar | LinkedIn

RESEARCH FOCUS

My research focuses on how language agents form abstractions, reason, and adapt continually. I aim to develop agentic systems capable of scientific exploration, robust to mode collapse and shallow chain-of-thought, with the goal of achieving structured and systematic reasoning.

EDUCATION

Ph.D. – McGill University & Mila - Québec AI Institute — Montreal, Canada
Advised by Prof. Blake A. Richards & Prof. Doina Precup — *2020 – Present*

M.Sc. – Concordia University — Montreal, Canada
Advised by Prof. Thomas Fevens — *2017 – 2019*
Thesis: Meta-Learning for Cancer Phenotype Prediction from Gene Expression Data [pdf]

B.Sc. – Shahid Beheshti University — Tehran, Iran
Advised by Prof. Mona Ghassemian — *2012 – 2016*

SELECTED RESEARCH PUBLICATIONS

Presenters are shwon in underline. Equal contribution is shwon in *.

1. **Language Agents Mirror Human Causal Reasoning Biases. How Can We Help Them Think Like Scientists?** Anthony GX Chen , Dongyan Lin*, **Mandana Samiei***, Doina Precup, Blake A. Richards, Rob Fergus, Kenneth Marino.
   Second conference on Language Modeling (COLM), 2025. [pdf]

2. **The Role of Schemas in Reinforcement Learning: Implications for Generalization.**
   Mandana Samiei, Doina Precup, Blake A. Richards. Conference on Reinforcement Learning and Decision Making (RLDM), 2025. [pdf]

3. **The Schema Spectrum: Explicit, Implicit, and Emergent Structures in AI and the Brain.**
   Mandana Samiei, Doina Precup, Blake A. Richards. Under review at Neuron, 2025. [pdf]

4. **AIF-GEN: Open-Source Platform and Synthetic Dataset Suite for Reinforcement Learning on Large Language Models.**
   Jacob Chmura*, Shahrad Mohammadzadeh*, Ivan Anokhin, Jacob-Junqi Tian, **Mandana Samiei**, Taz Scott-Talib, Irina Rish, Doina Precup, Reihaneh Rabbany, Nishanth Anand.
   Championing Open-Source Development in ML Workshop @ ICML, 2025. [OpenReview]

5. **Learning Schemas in Reinforcement Learning: bottleneck structure discovery.**
   Mandana Samiei, Doina Precup, Blake A. Richards.
   Under submission at Nature Communications, 2025.

6. **A conceptual analysis of continual learning objectives.**
   Giulia Lanzillotta, **Mandana Samiei**, Claire Vernade, Razvan Pascanu.
   Manuscript under submission to TMLR, 2025.

7. **Testing Causal Hypotheses through Hierarchical Reinforcement Learning.**
   Anthony GX Chen*, Dongyan Lin*, **Mandana Samiei***.
   Intrinsically-Motivated and Open-Ended Learning (IMOL) Workshop @ NeurIPS, 2024. [OpenReview]

8. **Mimicking Mammalian Navigation in Watermaze using Brain-Inspired Representations**
   Mandana Samiei*, Arna Ghosh*, Blake A. Richards Biological and Artificial Reinforcement Learning (BARL) Workshop at NeurIPS 2020. [Poster]

9. **Torchmeta: A Meta-Learning Library for PyTorch.**
Tristan Deleu, Tobias Würfl, **Mandana Samiei**, Joseph Paul Cohen, Yoshua Bengio.
PyTorch Developer Conference (PTDC), 2019. [arXiv]

Selected Invited Talks and Tutorials

The Role of Schemas in Reinforcement Learning Oct. 2025
Princeton Reinforcement Learning Lab *Remote*
- Invited research talk discussing schema representations in reinforcement learning and their implications for generalization.

Large Language Models – Tutorial [GitHub] Sept. 2025
Mediterranean Machine Learning Summer School (M2L) *Split, Croatia*
- Delivered a tutorial on the foundations and emerging research directions of large language models.

Towards Efficient Generalization in Continual RL using Episodic Memory Oct. 2021
Microsoft Research Summit 2021 *Remote*
- Invited talk on memory-augmented RL agents and their generalization efficiency, based on a collaboration work with Microsoft Research.

RL for Games – Tutorial [Notebook] Jul. 2021
Neuromatch Academy 2021 *Remote*
- Created an interactive tutorial on reinforcement learning for games, part of the Neuromatch Academy.

Awards and Recognition

Fonds de Recherche du Québec – Nature et Technologies (FRQNT) Doctoral Award 2022-2024
Women in AI Excellence Doctoral Scholarship, Mila 2021
UNIQUE (Unifying Neuroscience & AI) PhD Excellence Scholarship 2020
MITACS Accelerate Research Award 2019
National Merit Scholarship (top 0.2%) 2012

Selected Teaching & Mentorship

Tutor — Large Language Models, Mediterranean ML Summer School (M2L) 2025
Tutor — Introduction to ML, Eastern European Summer School (EEML) 2024
TA — Reinforcement Learning (COMP 579), McGill University 2022
TA — Fundamentals of Machine Learning (IFT 6390), University of Montreal 2021
TA — Intro to Robotics & Intelligent Systems (COMP 417), McGill University 2020

Service & Leadership

EDI Chair and Local Chair — Conference on Lifelong Learning Agents (CoLLAs) 2025, 2024
Board Member — Women in Machine Learning (WiML) 2022–Present
Reviewer — ICLR, NeurIPS, TMLR, CoLLAs, RLC, COLM 2022–2025
Organizer — ML Reproducibility Challenge (MLRC 2023) 2023
Organizer — Mila Neuro-AI Reading Group 2020-2023

Skills

Programming Languages, Libraries, & APIs: Python, Jax, Java, MATLAB, Javascript/CSS/HTML, Bash, Numpy, SciPy, Pandas, Jupyter, Git, SQL, Scikit-learn, TensorFlow, Slurm
Foundational Models, Fine-tuning and Inference: HuggingFace, vLLM, verl, Olama, OpenAI API
Spoken & Written: English (Fluent), Persian (native), French (Fluent), German (B1)

References

Prof. Doina Precup – McGill University, Mila, Google DeepMind dprecup@cs.mcgill.ca
Prof. Blake A. Richards – McGill University, Mila, Google Research blake.richards@mila.quebec
Prof. Thomas Fevens – Concordia University fevens@cs.concordia.ca
Prof. Mona Ghassemian - King's College London, Huawei mona.ghassemian@kcl.ac.uk