

Maxime Heuillet

Ph.D. candidate in Computer Science
Citizenship: French
3421, Rue Durocher #402
H2X 2C6, Montreal, QC, Canada

Phone: +1-438-792-1594
E-mail: maxime.heuillet.1@ulaval.ca
WWW: <https://maxheuillet.github.io>
Scholar: Google Scholar Profile

Education

Ph.D. in Computer Science 2021 - 2026
Université Laval and MILA - Québec AI Institute, Canada
Advisors: Prof. Audrey Durand, Dr. Ola Ahmad

M.Sc. in Statistics and Machine Learning 2017 - 2020
ENSAI - National School of Statistics and Data Analysis (french 'Grande École'), France.
- Admitted among 6 of 191 applicants; lowest admitted rank: 38.
- Research projects with: Prof Emmanuel Curris (Paris-Cité University), Prof. Julien Godet (Strasbourg University) and Prof. Romaric Gaudel (Inria & Rennes University)

Classe Préparatoire aux Grandes Écoles B/L (CPGE) 2015 - 2017
Aix Marseille University, joint program with Lycée Thiers, France.
- Two years intensive program to prepare for the French "Grandes Écoles" competitive exams
- **Majors:** mathematics, economics, social sciences

Scientific Achievements

My research experiences span reinforcement learning (**RL**) and its statistical foundations in **sequential decision making** (bandits, partial monitoring). I am also interested in the training and evaluation of large language models (**LLMs**) and their applications to recommender (**Recsys**), automated machine learning (**AutoML**) systems as well as emerging areas in **agentic AI** and **multi-agent** systems. I also worked on **robustness** in the **vision** domain.

Papers published in the main track of a conference

1. **M. Heuillet**, O. Ahmad, and A. Durand. *Randomized Confidence Bounds for Stochastic Partial Monitoring*, Proceedings of the International Conference on Machine Learning (2024). (tag: **Sequential decision making**)
2. **M. Heuillet**, O. Ahmad, and A. Durand. *Neural Active Learning Meets the Partial Monitoring Framework*, Proceedings of the Uncertainty in Artificial Intelligence conference (2024). (tag: **Sequential decision making**)

Pre-prints

1. **M. Heuillet**, Y. Cui, B. Chen, A. Durand, and P. Parthasarathi. *Nested-ReFT: Efficient Reinforcement Learning for Large Language Model Fine-Tuning via Off-Policy Rollouts*, To be submitted at ICML 2026. (tags: **RL**, **LLM**)
2. **M. Heuillet**, R. Bhagwatkar, J. Ngnawé, Y. Pequignot, A. Larouche, C. Gagné, I. Rish, O. Ahmad, and A. Durand. *A Guide to Robust Generalization: The Impact of Architecture, Pre-training, and Optimization Strategy*, To be submitted at TMLR 2026. (tags: **Robustness**, **Vision**)
3. B. Bonin, **M. Heuillet**, A. Durand. *LLM-as-a-Judge: Toward World Models for Slate Recommendation Systems*, To be submitted at ACL (2026) (tags: **LLM**, **Recsys**)
4. J. Ngnawé, **M. Heuillet**, S. Sahoo, Y. Pequignot, F. Precioso, C. Gagné, *Robust Fine-Tuning from Non-Robust Pre-trained Models: Mitigating Suboptimal Transfer With Epsilon-Scheduling*, Under review at ICLR (2026), (tags: **Robustness**, **Vision**)

Ongoing research projects

1. M. Heuillet, S. Peddijaru, B. Wang, S. Rajagopal. *CIPHER: A data science agent with test-time scaling to improve performance*, To be submitted at KDD 2026. (tags: **Agentic AI**, **LLM**, **AutoML**)

Workshop Posters

1. M. Heuillet, Y. Cui, B. Chen, A. Durand, and P. Parthasarathi. *Nested-ReFT: Efficient Reinforcement Learning for Large Language Model Fine-Tuning via Off-Policy Rollouts*, Accepted to Neurips 2025 First Workshop on Efficient Reasoning. (tags: **RL**, **LLM**)
2. J. Ngnawé, M. Heuillet, S. Sahoo, Y. Pequignot, F. Precioso, C. Gagné, *Robust Fine-Tuning from Non-Robust Pre-trained Models: Mitigating Suboptimal Transfer With Epsilon-Scheduling*, Accepted to Neurips workshop on reliable ML, (tags: **Robustness**, **Vision**)
3. M. Heuillet, R. Bhagwatkar, J. Ngnawé, Y. Pequignot, A. Larouche, C. Gagné, I. Rish, O. Ahmad, and A. Durand. *A Guide to Robust Generalization: The Impact of Architecture, Pre-training, and Optimization Strategy*, Accepted to Neurips workshop on reliable ML. (tags: **Robustness**, **Vision**)
4. J. Armand, T. C.-H. Lin, M. Heuillet, and A. Durand. *Multi-Agent Matrix Games with Individual Learners: How Exploration-Exploitation Strategies Impact the Emergence of Coordination*, Reinforcement Learning Conference 2025 Workshop on Coordination and Cooperation in Multi-Agent Reinforcement Learning (spotlight 🌟). (tags: **Sequential decision making**, **Multi-agent**)
5. M. Godbout, M. Heuillet, S. Chandra, R. Bhati, and A. Durand. *A Game-Theoretic Perspective on Risk-Sensitive Reinforcement Learning*. AAAI Workshop on Artificial Intelligence Safety a, 2022 (spotlight 🌟). (tags: **RL**, **Robustness**)
6. M. Heuillet, Benoit Debaque, and A. Durand. *Sequential AutoML: Bandits-driven Exploration using a Collaborative Filtering Representation*. 8th ICML Workshop on Automated Machine Learning, (2021). (tags: **AutoML**, **Sequential decision making**)
7. M. Traglia, M. Heuillet, and L. Weiss. *Male control individuals show higher burden of rare coding SNVs across diverse ancestries* In Proceedings of the American Society of Human Genetics Conference (2019) (tags: **Health**)

Professional Experience

Amazon Web Services (Seattle, USA) Applied Scientist Intern (L5) at Amazon Web Services	Fall 2025
Huawei Technologies - Noah's Ark Lab (Montréal, Canada) Associate Research Intern (part-time) Placement via Quantum Technologies Recruiting Inc.	Winter Summer 2025
Thales Research and Technology - CortAIx Lab (Montréal, Canada) Research Intern	Fall 2020 - Now
Université Laval (Québec, Canada) Teaching Assistant	2021, 2022
Center for Research in Economics and Statistics - CREST Lab (Rennes, France) Research Intern	Summer 2020
University of California San Francisco - Weiss Lab (Bay area, USA) Research intern	Summer 2019

Mentorship

Mentorship experiences part of work contracts (1 hour / week) or volunteering experiences.

Masters theses

- | | |
|---|----------------------|
| 1. Baptiste Bonin – Research M.Sc. at Université Laval
Recommending structured sequences of items | Fall 2024 - TBD |
| 2. Julien Armand – Research M.Sc. at Université Laval
Characterizing multi-agent learning dynamics (cooperation, competition) | Winter - Summer 2025 |

Internships

- | | |
|---|------------------------|
| 1. Vincent Longpré – Professional M.Sc. internship, at <i>Caisse de dépôt et placement du Québec</i>
"LLMs as a judge" techniques to assess quality of information extraction | Winter, Summer 2025 |
| 2. Donald Zvada – Professional M.Sc. internship at <i>African Inst. of Math. Sci. in South Africa</i>
Can large language models solve sequential decision making problems? | Winter, Summer 2025 |
| 3. Frida-Cecilia Acosta Parenteau – Professional M.Sc. internship, at <i>SmartOne.AI</i> startup
Predicting the cost of human data annotations | Summer, Fall 2024 |
| 4. Myyank Shukla – Professional M.Sc. internship, at <i>Royal Bank of Canada</i>
Monitoring for anomaly detection | Fall 2023, Winter 2024 |
| 5. Lorenzo Mathieu – Professional M.Sc. internship at <i>Université Laval</i>
Transfer learning from synthetic to real world data | Summer 2023 |
| 6. Thierry Blais – B.Sc. internship at <i>Université Laval</i>
Neural bandits | Summer, Fall 2022 |

Awards and scholarships

- | | |
|---|-------------|
| Mitacs Accelerate, Doctoral level Scholarship
135.000\$ to complete my Ph.D. thesis
<i>Partner organization:</i> Thales Research and Technology (CortAIx lab) | 2022 - 2026 |
| Bourse mobilité étudiante (CRDM)
2.000\$ to attend the Neurips conference in San Diego (CA) | 2025 |
| IVADO Travel grant
3.000\$ to attend a scientific workshop at the University of California in Berkeley (USA)
Topic: Future of Language Models and Transformers | 2025 |
| Alliance Canada Computing Resources Award
Led the application to obtain computing resources for Pr. Durand's group.
2024 - 2025: 11.453\$ for "Online Learning: Bridging Theory and Practice"
2025 - 2026: 21.007\$ for the award renewal | 2024 - 2026 |
| UAI 2024 Conference Scholarship
500\$ fee waiver to attend the UAI 2024 conference. | 2024 |
| Université Laval Mobility Scholarship
500\$ to present at a scientific conference outside Québec | 2024 |
| Scholarship of Excellence in AI and Data Science
7.000\$ from Institute Intelligence and Data (Université Laval) | 2022 |
| Université Laval Admission Grant (2.000\$) | 2021 |
| Mitacs Accelerate, Master's level Scholarship
15.000\$ for "Recommender Systems for Automated Machine Learning"
<i>Partner organization:</i> Thales Research and Technology (CortAIx lab) | 2020 |
| Bourse Jeune à Travers le Monde (JTM)
270\$ for an internship abroad in Montréal, Québec | 2020 |
| Bourse Jeune à l'International (JALI)
4.000\$ for an internship abroad in San Francisco, California | 2019 |

Community Service

Peer-review contributions

Conference (main-track) Reviewer: ICML 2024, Canadian AI Conference 2024, AAAI 2025

Workshop Reviewer:

Reinforcement Learning for Real Life (NeurIPS 2022 workshop)

Montreal AI Symposium 2022, Decision Awareness in Reinforcement Learning (ICML 2022 workshop)

Pac-Bayes Meets Interactive Learning (ICML 2023 workshop)

Goal Conditioned Reinforcement Learning (Neurips 2023 workshop)

Explainable AI for Speech and Audio (ICASSP 2024 workshop)

Workshop on Training Agents with Foundation Models (RLC 2024)

Montreal AI Symposium 2024

Program Chair Pac-Bayes Meets Interactive Learning (ICML 2023 workshop),

Admission Committee:

Reviewed international student applications for MILA Ph.D. and M.Sc. programs

Designed fit, technical, and ML interviews for admissions into Université Laval Ph.D. program

Scientific Events Organization

Mila 2024 Workshop on Reinforcement Learning

2023

One day event featuring distinguished speakers and panelists in the field of reinforcement learning.

ICML 2023 Workshop "Pac-Bayes Meets Interactive Learning" ¹

2023

One day event featuring distinguished speakers and panelists in the field of Pac-Bayes and interactive learning.

AI tutorials - Université Laval bootcamp 2022 ²

2022

Two-days event with AI tutorials addressed to a non-expert audience.

ENSAI International alumni day (North America alumni network)

2021

A day to connect current ENSAI students with international alumnis of North America.

ML Robustness reading group ³

2022

Created a bi-monthly reading group for students interested in the robustness and security of AI

Talks

1. *Improving cost efficiency across the life cycle of machine learning systems*

"Carrefour DEEL France - Québec", a transatlantic scientific seminar (2025)

"Thales Palaiseau Days", a Thales Group internal conference (2024)

2. *Introduction to Reinforcement Learning*

"AI Bootcamp for students", a department workshop at Université Laval (2022)

3. *Sequential Pipeline Optimization: Bandits-driven Exploration using a Collaborative Filtering Representation*

"International Workshop on Automation in Machine Learning" at KDD Conference (2021)

Montréal AI Symposium at Mila (2021)

Departmental seminars at Université Laval (2022)

¹<https://bguedj.github.io/icml2023-workshop/>

²<https://sites.google.com/view/bootcamp-iid-2022/accueil?authuser=1>

³<https://sites.google.com/view/adversarial-robust-ml-reading/accueil?authuser=1>