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

Education

Ph.D. in Computer Science 2021 - 2026
Université Laval and *MILA - Québec AI Institue*, 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 **sequential decision making** (bandits, partial monitoring), the training and evaluation of large language models (**LLMs**) and applications to recommender systems (**Recsys**) and automated machine learning (**AutoML**), as well as questions of **robustness** across **vision**, **health**, and emerging areas in **agentic AI** and **multi-agent** systems.

Articles in Preparation

1. A data science agent to promote accessibility of machine learning technology (tags: **Agentic AI**, **LLM**, **AutoML**)
2. How can LLMs best solve sequential decision making problems? (tags: **LLM**, **Sequential decision making**)
3. Leveraging LLM judges for recommendation systems of sequences of items (tags: **LLM**, **Recsys**)

Submitted Articles

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*, Submitted to AAAI 2026 (Main Track). (tags: **RL**, **LLM**)
2. **M. Heuillet**, R. Bhagwatkar, J. Ngnawe, 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*, Submitted to the NeurIPS Datasets and Benchmarks Track (2025). (tags: **Robustness**, **Vision**)
3. Jonas Ngnawé, **Maxime Heuillet**, Sabyasachi Sahoo, Yann Pequignot, Frédéric Precioso, Christian Gagné, *Robust Fine-Tuning from Non-Robust Pretrained Models: Mitigating Suboptimal Transfer With Epsilon-Scheduling*, Submitted to Neurips workshop on reliable ML, (tags: **Robustness**, **Vision**)
4. M. Godbout, **M. Heuillet**, and A. Durand. *An Adversarial Perspective on Risk-Sensitive Reinforcement Learning*, Submitted to AAAI 2026 (Main Track). (tags: **RL**, **Robustness**)

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**)

Workshop Posters

1. 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 Learn- ing (**spotlight** 💫). (tags: **Sequential decision making**, **Multi-agent**)
2. M. Godbout, **M. Heuillet**, S. Chandra, R. Bhati, and A. Durand. *A Game-Theoretic Perspective on Risk-Sensitive Reinforcement Learning*. AAAIs Workshop on Artificial Intelligence Safety a, 2022 (**spotlight** 💫). (tags: **RL**, **Robustness**)
3. **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**)
4. 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 (Seattle, USA) Fall 2025
Applied Scientist Intern

Huawei Technologies - Noahs Ark Lab (Montréal, Canada) Winter Summer 2025
Associate Research Intern (part-time)
Placement via Quantum Technologies Recruiting Inc.

Thales Research and Technology - CortAIx Lab (Montréal, Canada) Fall 2020 - Now
Research Intern

Université Laval (Québec, Canada) 2021, 2022
Teaching Assistant

Center for Research in Economics and Statistics - CREST Lab (Rennes, France) Summer 2020
Research Intern

University of California San Francisco - Weiss Lab (Bay area, USA) Summer 2019
Research intern

Mentorship

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

Masters theses

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

Internships

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

Awards and scholarships

Mitacs Accelerate, Doctoral level Scholarship 2022 - 2026
135.000$ to complete my Ph.D. thesis
Partner organization: Thales Research and Technology (CortAIx lab)

IVADO Travel grant 2025
3.000$ to attend a scientific workshop at the University of California in Berkeley (USA)
Topic: Future of Language Models and Transformers

Alliance Canada Computing Resources Award 2024 - 2026
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

UAI 2024 Conference Scholarship 2024
500$ fee waiver to attend the UAI 2024 conference.

Université Laval Mobility Scholarship 2024
500$ to present at a scientific conference outside Québec

Scholarship of Excellence in AI and Data Science 2022
7.000$ from Institute Intelligence and Data (Université Laval)

Université Laval Admission Grant (2.000$) 2021

Mitacs Accelerate, Master's level Scholarship 2020
15.000$ for "Recommender Systems for Automated Machine Learning"
Partner organization: Thales Research and Technology (CortAIx lab)

Bourse Jeune a Travers le Monde (JTM) 2020
270$ for an internship abroad in Montréal, Québec

Bourse Jeune a l'International (JALI) 2019
4.000$ for an internship abroad in San Francisco, California

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 adressed 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