ANTOINE MOULIN

PhD in Machine Learning

Homepage ◇ Google Scholar ◇ Email: antoine.moulin@berkeley.edu

RESEARCH INTERESTS

Reinforcement learning (efficient exploration with function approximation, adversarial environments, distribution shift), post-training of language models.

EDUCATION

Ph.D in Machine Learning Dec 2021 – Oct 2026
Universitat Pompeu Fabra Barcelona, Spain

- The overarching topic is to identify and exploit structural properties of Markov decision processes that make large-scale reinforcement learning (RL) and imitation learning (IL) statistically and computationally feasible.
- ELLIS visiting student at the Gatsby Unit, University College London. Advisors: Gergely Neu, Arthur Gretton.

Master of Science in Machine Learning Sep 2019 – Sep 2020
ENS Paris-Saclay Paris, France

- Grade: 17.34/20.00. Courses: RL, Computer Vision, Graphical Models, Convex Optimization, Deep Learning, Neurosciences, Kernel Methods, Sequential Learning, Multi-scale models, Biostatistics.

Engineering Degree Sep 2017 – Sep 2020
Telecom Paris Paris, France

- GPA: 4.00/4.00. Equivalent to a Masters, top 5% cohort. Some courses: Measure & Probability Theory, Functional Analysis, Signal Processing, Statistics, Machine Learning, Optimization, Algorithm Design.

Preparatory Classes Sep 2015 – Jul 2017
Lycée Janson-de-Sailly Paris, France

- GPA: 3.96/4.00. Post-secondary two-year program in Advanced Mathematics and Physics.

EXPERIENCE

Postdoctoral Fellow in Machine Learning Nov 2026 –
UC Berkeley Berkeley, CA

- Postdoc at the Simons Institute at UC Berkeley, hosted by Peter Bartlett.

Student Researcher Jun 2024 – Dec 2024
Google DeepMind London, UK

- Developed finite-time predictive interval methods for ensemble-based uncertainty quantification under general function classes and approximations thereof. Advisor: Ilja Kuzborskij.

Research Engineer Feb 2021 – Apr 2021
Research Intern May 2020 – Oct 2020
INRIA Lille, France

- Implemented extensions of actor-critic methods (PPO, TRPO) to improve exploration, worked on reproducibility issues in RL. Advisors: Yannis Flet-Berliac, Philippe Preux.

Research Intern Summer 2019
Aalto University Helsinki, Finland

- Implemented an efficient extension of a sparse kernel CCA method. Advisor: Juho Rousu.

Analyst Intern Summer 2018
Crédit Agricole Caen, France

PUBLICATIONS & PREPRINTS (* Equal contribution)

When Does On-Policy Interaction Help? Representational Tradeoffs in Value-Based IL [arXiv]
Luca Viano*, **Antoine Moulin***, Audrey Huang, Volkan Cevher, Philip Amortila, and Dylan J. Foster.
(NeurIPS 2026) 40th Annual Conference on Neural Information Processing Systems.
Contributed talk at the 19th European Workshop on Reinforcement Learning (EWRL 2026).

Outcome-Aware Spectral Feature Learning for Instrumental Variable Regression [arXiv]
Dimitri Meunier, Jakub Wornbard, Vladimir R Kostic, **Antoine Moulin**, Alek Frölich, Karim Lounici, Massimiliano Pontil, and Arthur Gretton.
(ICML 2026) 43rd International Conference on Machine Learning.
Best poster award at the Citadel PhD Summit 2026 (London).

QEDBench: Quantifying the Alignment Gap in Automated Evaluation of University-Level Mathematical Proofs [arXiv]
Santiago Gonzalez, Alireza Amiri, Peter Ye, Edward Zhang, . . . , **Antoine Moulin**, . . . , Quanquan C. Liu.
(ICML 2026) 43rd International Conference on Machine Learning.

A benchmark of expert-level academic questions to assess AI capabilities [Open access] [arXiv]
Center for AI Safety, Scale AI, and **HLE Contributors Consortium**.
Nature, 2026.

Inverse Q-Learning Done Right: Offline Imitation Learning in Q^π-Realizable MDPs [arXiv]
Antoine Moulin, Gergely Neu, and Luca Viano (α-β order).
(NeurIPS 2025) 39th Annual Conference on Neural Information Processing Systems.

Demystifying Spectral Feature Learning for Instrumental Variable Regression [arXiv]
Dimitri Meunier, **Antoine Moulin**, Jakub Wornbard, Vladimir R Kostic, and Arthur Gretton.
(NeurIPS 2025) 39th Annual Conference on Neural Information Processing Systems.

When Lower-Order Terms Dominate: Improved Loss-Range Adaptivity for Expert Algorithms [arXiv]
Antoine Moulin*, Emmanuel Esposito*, and Dirk van der Hoeven.
(NeurIPS 2025) 39th Annual Conference on Neural Information Processing Systems.

Optimistically Optimistic Exploration for Provably Efficient Infinite-Horizon RL and IL [arXiv]
Antoine Moulin, Gergely Neu, and Luca Viano (α-β order).
(COLT 2025) 38th Annual Conference on Learning Theory.
Contributed talk at the 18th European Workshop on Reinforcement Learning (EWRL 2025).

Spectral Representation for Causal Estimation with Hidden Confounders [arXiv]
Haotian Sun*, **Antoine Moulin***, Tongzheng Ren, Arthur Gretton, and Bo Dai.
(AISTATS 2025) 28th International Conference on Artificial Intelligence and Statistics.

Optimistic Planning by Regularized Dynamic Programming [arXiv]
Antoine Moulin, and Gergely Neu (α-β order).
(ICML 2023) 40th International Conference on Machine Learning.

TALKS

Value-Based Methods for Imitation Learning — Sep 2026
Netflix, MLIR team — Remote

Efficient exploration in linear Markov decision processes — Jan 2026
University of Oxford — Oxford, UK

Tutorial on Imitation Learning — Jan 2026
University of Oxford — Oxford, UK

Efficient exploration in linear Markov decision processes Nov 2025
Isaac Newton Institute Cambridge, UK

Inverse Q-Learning for Offline Imitation Learning Sep 2025
Università degli Studi di Milano Milan, Italy

Learning in Adversarial Linear MDPs Mar 2024
University of Tokyo Tokyo, Japan

Optimistic Planning by Regularized Dynamic Programming Aug 2023
Princeton University Princeton, NJ

Optimistic Planning by Regularized Dynamic Programming Jul 2023
Stanford University Stanford, CA

Infinite-Horizon RL with Linear Function Approximation Mar 2023
Universitat Pompeu Fabra Barcelona, Spain

Primal-dual Methods for RL Sep 2022
Gatsby Unit, University College London London, UK

Introduction to JAX Sep 2021
ELLIS Doctoral Symposium Tubingen, Germany

TEACHING

Machine Learning	Universitat Pompeu Fabra	2022 – 2023
Algorithm Design	Universitat Pompeu Fabra	2022 – 2024
Computer Science	IPESUP, bootcamp for high-school students	2023
Probabilities	Universitat Pompeu Fabra	2022
Mathematics Examiner	Lycée Janson-de-Sailly, preparatory classes	2018 – 2019

ACADEMIC ACTIVITIES

- Organizer of the NeurIPS 2025 workshop *Aligning Reinforcement Learning Experimentalists and Theorists*.
- Organizer of the RL theory virtual seminars (YouTube channel).
- Organizer of the ICML 2024 workshop *Aligning Reinforcement Learning Experimentalists and Theorists*.
- Organizer of the Reinforcement Learning Summer School 2023 (RLSS 2023).
- Reviewer: EWRL (2023-25), ARLET (2024-25), NeurIPS (2025), JMLR (2025), ICLR (2026), AISTATS (2026).
- Chair of the ELLIS RL reading group (2021-2022).

SKILLS

Programming	Python (JAX, PyTorch), MATLAB, C++, Java, SQL, HTML.
Languages	French (native), English (fluent), Spanish, Italian (intermediate).
Sports	Climbing, (beach) volleyball, ski, table tennis.

SCHOLARSHIPS & GRANTS

IMSI Long Program Scholarship	2026
ELISE/ELSA mobility grants (ELLIS)	2022, 2024, 2025
Doctoral scholarship, Universitat Pompeu Fabra	2021
Declined: Doctoral scholarships at U. of Cambridge, U. College London	2021