Alessandro Montenegro | Ph.D. Candidate

Milan, Italy
+39 3933903971 • alessandro.montenegro@polimi.it • alessandro-montenegro
MontenegroAlessandro • MontenegroAlessandro

Short Bio

Alessandro Montenegro is a Ph.D. Student in Information Technology at Politecnico di Milano (expected graduation: December 2026), working under the supervision of Profs. Alberto Maria Metelli and Matteo Papini. He holds a Master's degree in Computer Science and Engineering from Politecnico di Milano (2023) and a Bachelor's degree from University of Rome Tor Vergata (2020), both achieved with honors. Previously a research assistant at the AIRLab, his current research focuses on Reinforcement Learning, specifically investigating the sample efficiency of policy-based methods in real-world-like settings. His work bridges theoretical foundations with practical applications, including industrial projects and robot learning.

Education

Ph.D. in Information Technology — **Milan, Italy**
Politecnico di Milano — *Sep. 2023 – Dec. 2026 (3 Years)*
Topic: Bridging the Theory-Practice Gap in Policy-based Methods of Reinforcement Learning.
Supervisor: Prof. Alberto Maria Metelli
Co-supervisors: Prof. Matteo Papini, Dr. Marco Mussi
Research highlights: Published 3 papers at ICML (one of which being a spotlight one, top 3.5%), 1 paper at NeurIPS, and 1 paper at RSS (all A* venues). Research spans the following main threads: (i) policy gradient convergence theory [C5] with dynamic stochasticity [C3], (ii) policy gradient methods for constrained MDPs [C4]-[R2], (iii) sample-efficient trajectory reuse in policy gradient algorithms [C1]-[R1]-[P1].

M.Sc. in Computer Science and Engineering — **Milan, Italy**
Politecnico di Milano — *Sep. 2020 – May 2023 (2.5 years)*
Main focus: Artificial Intelligence and Machine Learning
Master thesis: Best Model Selection via Stochastic Rising Bandits
Supervisor: Prof. Alberto Maria Metelli
Final Mark: 110/110 cum Laude
Highlights: Thesis published at ICML as a spotlight paper [C6] (top 3.5%).

B.Sc. in Computer Science and Engineering — **Rome, Italy**
Università degli Studi di Roma Tor Vergata — *Sep. 2017 – Oct. 2020 (3 Years)*
Main focus: Robotics and Automation
Dissertation: *Projects Portfolio*
Supervisor: Prof. Francesco Lo Presti
Final Mark: 110/110 cum Laude

High School Diploma — **Ciampino (RM), Italy**
Liceo Scientifico Statale Vito Volterra — *Sep. 2012 – Jul. 2017 (5 Years)*
Dissertation: *The Imitation Game: History and Frontiers of Artificial Intelligence*
Final Mark: 100/100 cum Laude

Research and Teaching Experience

DFKI and TU Darmstadt — **Darmstadt, Germany**
Visiting Research Fellow — *Sep. 2025 – Dec. 2025 (3 Months)*
Visiting PhD Student at DFKI (German Research Center for Artificial Intelligence).
Topic: Accurate locomotion and foothold tracking for humanoid robots.
Project Responsibles: Prof. Puze Liu, Prof. Jan Peters
Highlight: Produced a paper titled "Mind Your Steps: A General Learning Framework for Accurate Humanoid Foothold Tracking" [C2] whose focus is to investigate novel lightweight paradigms for achieving accurate locomotion with humanoids via Reinforcement Learning.

Politecnico di Milano — **Milan, Italy**
Teaching Assistant — *Sep. 2024 – Aug. 2025 (1 Year)*

Course Name: "Informatica A" (Computer Science)
Details: 10 C.F.U. course issued in Italian for first year Mathematical Engineering students
Course Responsible: Prof. Giacomo Boracchi
Website: `https://montenegroalessandro.github.io/InfoA2425`
Activities: Exercise lecturer, written exam preparation and correction, oral exam execution
Duration: 38h
Student Ranking: 3.3/4 (Avg. University 3.1/4)
Student Attendance: ≈ 150

Politecnico di Milano **Milan, Italy**
Doctoral Student (via Assegno di Ricerca) *Sep. 2023 – Dec. 2026 (3 Years)*
Artificial Intelligence and Robotics Lab (AIRLab)
Supervisor: Prof. Alberto Maria Metelli
Co-supervisors: Prof. Matteo Papini, Dr. Marco Mussi
Research Highlights:

- **Policy gradient convergence**: [C5]-[C3]-[W1]
- **Policy gradients for constrained MDPs**: [C4]-[R2]
- **Sample-efficient trajectory reuse**: [C1]-[R1]-[P1]-[W2]
- **Humanoid Locomotion via Reinforcement Learning**: [C2]

Others: Supervised **3 M.Sc. theses**, supervised **2 research assistants**, **sponsorship chair** and **staff manager** of RLSS (Reinforcement Learning Summer School) 2026 (`https://rlsummerschool.com/`).

Politecnico di Milano **Milan, Italy**
Research Assistant (Assegno di Ricerca) *Jun. 2023 – Sep. 2023 (3 Months)*
Artificial Intelligence and Robotics Lab (AIRLab)
Supervisor: Prof. Alberto Maria Metelli
Activity: Participated to an industrial research project with Leonardo S.p.A, participated to the "AirHockey Challenge 2023", supervised 1 M.Sc. Thesis.

Languages

Italian: Mother Tongue

English: Excellent (IELTS 7)

Research Projects

Competitive Research Projects

FAIR (Future of Artificial Intelligence Research) - Spoke 4: Adaptive AI **Milan, Italy**
Extended Partnership - National Recovery and Resilience Plan (PNRR) *Jun. 2023 – Mar. 2026*
Focus: Artificial intelligence: foundational aspects - Adaptive AI
Spoke coordinator: Politecnico di Milano (Prof. Nicola Gatti)
Funding institution: Italian Ministry of University and Research (MUR)
Main Activity: investigation of reinforcement learning techniques to bridge the theory-practice gap in policy-based methodologies. Specifically, the research activity realized either in establishing theoretical frameworks for explaining common practices employed in deep RL methods, or in developing new algorithms with strong convergence guarantees and empirical performance addressing real-world challenges. This research activities led to the following research outcomes: [C5]-[C3]-[C4]-[C1]-[W2]-[W1]-[R1]-[R2]. *Work done in parallel with the PhD.*

Competitive Research Projects for Computational Time

PGCLaSI - Policy Gradients for Control of Large-Scale Industrial Plants **Milan, Italy**
CINECA Supercomputing Centre - Class C Project *Oct. 2023 – Jul. 2024*
Principal Investigator: Prof. Alberto Maria Metelli
Topic: Development of efficient policy gradient algorithms for large-scale industrial plants
Grant: 80k GPU hours
Main Activity: development of policy-based reinforcement learning methods for constrained settings under continuous state-action spaces and simultaneously considering multiple (risk-based) constraints. This research activity led to the following research outcomes: [C4]-[R2]. *Work done in parallel with the PhD.*

Industrial-Funded Research Projects

Collaboration with Leonardo S.p.A. **Milan, Italy**
Applied Machine Learning *Jun. 2023 – Sep. 2023*

Role: Applied research scientist
Co-principal investigators: Proff. Marcello Restelli and Nicola Gatti

Collaboration with ENI S.p.A. **Milan, Italy**
Applied Machine Learning *Jun. 2026 – Now*
Role: Project manager
Co-principal investigators: Proff. Marcello Restelli, Alberto Maria Metelli, and Marco Mussi

Collaboration with Leonardo S.p.A. **Milan, Italy**
Applied Machine Learning *Sep. 2026 – Now*
Role: Project manager
Co-principal investigators: Proff. Marcello Restelli, Alberto Maria Metelli, and Marco Mussi

Challenges

AirHockey Challenge 2023 **Milan, Italy**
TU Darmstadt *Feb. 2023 – Nov. 2023*
Final Placement: 5th out of 46 participants
Related output: M.Sc. thesis supervision of Francesco Minnucci (December 2023).
Main Activity: developed reinforcement learning-based controllers for a general-purpose robot arm to play air hockey, combining rule-based planning with learned policies.

Research Software

- Developed **MagicRL** and **MagicRL3**, a research framework implementing the RL interaction protocol and a suite of policy gradient algorithms. Used as the backbone for [C4]-[C5]-[C3]-[C1]-[W1]-[W2]-[R2]. Code:
  1. `https://github.com/MontenegroAlessandro/MagicRL`
  2. `https://github.com/MontenegroAlessandro/MagicRL3`
- Extended a JAX-based humanoid locomotion framework for policy training and reliable hardware deployment. Used in [C2]; project page: `https://montenegroalessandro.github.io/mind-your-steps/`. Code: `https://github.com/MontenegroAlessandro/loco-mujoco`.

Awards and Grants

- Gold Reviewer, top 25% (2026) - ICML (International Conference on Machine Learning)
- Top Reviewer (2025) - NeurIPS (Neural Information and Processing Systems) conference
- Travel grant (2024) - NeurIPS (Neural Information and Processing Systems) conference
- Exemption from university fees due to high GPA (2020 - 2023) - Politecnico di Milano
- Winner of *Assegno di Incentivazione* Scholarship (2020) - MIUR
- Exemption from university fees due to high GPA (2017 - 2020) - Università degli Studi di Roma Tor Vergata
- Enrollment to the *Albo Nazionale Eccellenze* (2017) - INDIRE - National register of the best Italian high-school graduates

Publications

Conferences

[C1] Alessandro Montenegro, Federico Mansutti, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Reusing Trajectories in Policy Gradients Enables Fast Convergence. 2026. ICML, International Conference on Machine Learning (A* Core Ranking). Acceptance Rate: 26.6%.

[C2] Alessandro Montenegro, Shihao Li, Puze Liu, Alberto Maria Metelli, Jan Peters. Mind Your Steps: A General Learning Framework for Accurate Humanoid Foothold Tracking. 2026. RSS, Robotics: Science and Systems (A* Core Ranking). Acceptance Rate: 29.7%.

[C3] Alessandro Montenegro, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Convergence Analysis of Policy Gradient Methods with Dynamic Stochasticity. 2025. ICML, International Conference on Machine Learning (A* Core Ranking). Acceptance Rate: 26.9%.

[C4] Alessandro Montenegro, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Last-Iterate Global Convergence of Policy Gradients for Constrained Reinforcement Learning. 2024. NeurIPS, Advances in Neural Information Processing Systems (A* Core Ranking). Acceptance Rate: 25.8%.

[C5] Alessandro Montenegro, Marco Mussi, Alberto Maria Metelli, Matteo Papini. Learning Optimal Deterministic Policies with Stochastic Policy Gradients. 2024. ICML, International Conference on Machine Learning (A* Core Ranking). Acceptance Rate: 27.5%. Spotlight paper: top 3.5%

[C6] Marco Mussi, Alessandro Montenegro, Francesco Trovò, Marcello Restelli and Alberto Maria Metelli. Best Arm Identification for Stochastic Rising Bandits. 2024. ICML, International Conference on Machine Learning (A* Core Ranking). Acceptance Rate: 27.5%. Spotlight paper: top 3.5%.

Workshops

[W1] Alessandro Montenegro, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Convergence Analysis of Policy Gradient Methods with Dynamic Stochasticity. 2025. ELLIS Unconference at EurIPS.

[W2] Alessandro Montenegro, Federico Mansutti, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Reusing Trajectories in Policy Gradients Enables Fast Convergence. 2025. Salon de Refusés at EurIPS.

[W3] Alessandro Montenegro, Marco Mussi, Francesco Trovó, Marcello Restelli and Alberto Maria Metelli. Stochastic Rising Bandits: A Best Arm Identification Approach. European Workshop on Reinforcement Learning (EWRL). 2023.

[W4] Alessandro Montenegro, Marco Mussi, Francesco Trovó, Marcello Restelli and Alberto Maria Metelli. A Best Arm Identification Approach for Stochastic Rising Bandits. Workshop on New Frontiers in Learning, Control, and Dynamical Systems at International Conference of Machine Learning (ICML). 2023.

Under Review

[R1] Alessandro Montenegro, Riccardo Venturelli, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Reusing Samples in Proximal Policy Optimization: When and How Does It Help?. 2026.

[R2] Alessandro Montenegro, Leonardo Cesani, Marco Mussi, Matteo Papini, Alberto Maria Metelli. Learning Deterministic Policies with Policy Gradients in Constrained Markov Decision Processes. 2025.

In Preparation

[P1] Fine-tuning of LLMs via Retrospective Policy Gradient Methods. *(Applying trajectory reuse ideas from [C1] to RLHF-style LLM fine-tuning.)*

[P2] Optimal Reward Discovery in Robotics Applications via Human Feedback.

Master's Students and Research Assistants Supervision

- Caterina Motti (Research Assistant), industrial-funded research projects. Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Alessandro Montenegro, Marco Mussi, Marcello Restelli.
- Maria Cristina Centorame (Research Assistant), industrial-funded research projects. Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Alessandro Montenegro, Marco Mussi, Marcello Restelli.
- Riccardo Venturelli (Research Assistant), "Reusing Past Samples in Deep On-Policy Reinforcement Learning Methods". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Alessandro Montenegro, Marco Mussi, Matteo Papini. *Work led to [R1].*
- Federico Mansutti (M.Sc. Student), "Trajectory Reuse in Policy Gradients". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Alessandro Montenegro, Marco Mussi, Matteo Papini, Brian D. Ziebart. M.Sc. in Computer Science and Engineering, July 2025. *Work led to [C1] and [W2].*
- Leonardo Cesani (M.Sc. Student), "Learning Deterministic Policies in Constrained Markov Decision Processes with Policy Gradients". Supervisor: Prof. Matteo Papini. Co-supervisors: Alberto Maria Metelli, Alessandro Montenegro, and Marco Mussi. M.Sc. in Computer Science and Engineering, April 2025. *Work led to [R2].*
- Francesco Minnucci (M.Sc. Student), "Applying rule-based controllers and reinforcement learning to control a general purpose robot: the Air Hockey challenge case". Supervisor: Prof. Marcello Restelli. Co-supervisors: Amarildo Likmeta and Alessandro Montenegro. M.Sc. in Computer Science and Engineering, December 2023.

Other Academic Activities

Reviewer Activities

- International Conference on Machine Learning (ICML) 2024, 2025, 2026
- Conference on Neural Information and Processing Systems (NeurIPS) 2025, 2026
- Aligning Reinforcement Learning Experimentalists and Theorists (ARLET) Workshop 2024 (ICML), 2025 (NeurIPS)
- Exploration in AI Today (EXAIT) Workshop 2025 (ICML)
- IEEE Robotics and Automation Letters Journal
- Robotics and Autonomous Systems Journal
- PLOS One Journal
- Reviewer for the ELLIS PhD's Applications 2025
- Transactions on Machine Learning Research (TMLR) Journal

Participation to International Conferences and Workshops

- RSS 2026, Robotics: Science and Systems.
  Sydney, Australia. July 2026.
- ICML 2026, International Conference on Machine Learning.
  Seoul, South Korea. July 2026.
- ICML 2025, International Conference on Machine Learning.
  Vancouver, Canada. July 2025.
- NeurIPS 2024, Conference on Neural Information Processing Systems.
  Vancouver, Canada. December 2024.
- ICML 2024, International Conference on Machine Learning.
  Vienna, Austria. July 2024.
- EWRL 2023, European Workshop on Reinforcement Learning.
  Brussels, Belgium. September 2023.

Participation to Summer Schools

- DLRL 2024, Deep Learning and Reinforcement Learning.
  One poster presented.
  Toronto, Canada. July 2024.

Seminars and Appearance

- *Policy Gradients: Convergence and Fast Variants*. Seminar delivered at TU Darmstadt. October 31, 2025.
- *Model-based Reinforcement Learning*. Seminar delivered at DEIB, Politecnico di Milano. June 13, 2025.
- *Learning Optimal Deterministic Policies with Stochastic Policy Gradients* (ICML 2024) is on the Reading List of RL Theory Seminars. `https://sites.google.com/view/rltheoryseminars/reading-list`.
- *Learning Optimal Deterministic Policies with Stochastic Policy Gradients* (ICML 2024) received a dedicated episode at *Root Access* podcast. `https://youtu.be/leDS8Ypqt_s?si=uvDLXP5OXnbE22Lj`.

Other Activities

- Sponsorship Chair and Staff Manager for the Reinforcement Learning Summer School (RLSS) 2026 (`https://rlsummerschool.com/`)
- Voluntary activities for earthquake victims (2017)

- Voluntary activities for AIL (Associazione italiana contro le leucemie) (2015 – 2017)
- High School Students Representative (Sep. 2015 - Jul. 2017)

Last updated: September 9, 2026