

Mirco Mutti

Postdoc at the Technion

✉ muttimirco@gmail.com

Research Interest

I am interested in all the facets of *reinforcement learning*. My aim is to advance theoretical understanding that can lead to successful application of reinforcement learning in the real world.

Education

- 2023 **Ph.D. in Data Science and Computation**, Politecnico di Milano and Università di Bologna.
Thesis: Unsupervised Reinforcement Learning via State Entropy Maximization.
Advisor: Marcello Restelli (*Reviewers:* Mohammad Gheshlaghi Azar, Lerrel Pinto).
Grade: Excellent.
- 2018 **M.Sc. in Computer Science and Engineering**, Politecnico di Milano.
Thesis: Configurable Markov Decision Processes.
Advisor: Marcello Restelli.
Grade: 110/110 cum laude.
- 2015 **B.Sc. in Engineering of Computing Systems**, Politecnico di Milano.
Grade: 110/110 cum laude.

Positions

- 2023-Present **Postdoc**, Technion – Israel Institute of Technology.
Research Interest: Reinforcement learning from theory to practice.
Advisor: Aviv Tamar.
- 2018 **Research fellow**, Politecnico di Milano.
Research: Predictive models of disruptive events during drilling commissioned by ENI S.p.a.
Advisor: Marcello Restelli.

Teaching Assistant

- Spring 2022 **Machine learning**, M.Sc. at Politecnico di Milano, 10 hrs of exercise sessions, \approx 400 students.
- Fall 2021 **Informatics B**, B.Sc. at Politecnico di Milano, 26 hrs exercise session, \approx 200 students
- Spring 2021 **Machine learning**, M.Sc. at Politecnico di Milano, 18 hrs of exercise sessions, \approx 400 students.
- Fall 2020 **Informatics B**, B.Sc. at Politecnico di Milano, 36 hrs exercise session, \approx 200 students.
- Spring 2020 **Machine learning**, M.Sc. at Politecnico di Milano, 10 hrs of exercise sessions, \approx 400 students.
- Fall 2019 **Informatics B**, B.Sc. at Politecnico di Milano, 26 hrs of exercise sessions, \approx 200 students.
- Fall 2018 **Informatics B**, B.Sc. at Politecnico di Milano, 26 hrs of exercise sessions, \approx 200 students.

Honors

- 2023 **Honorable mention** for the best PhD thesis on artificial intelligence granted by the Italian Association for Artificial Intelligence (AIxIA).
- 2023 **Notable paper** at AISTATS 2023.
- 2022 **Outstanding paper award** at ICML 2022.
- 2019 **Thesis award** for the best Master's thesis on artificial intelligence granted by the Italian Association for Artificial Intelligence (AIxIA).
- 2018 **Ph.D. scholarship** for outstanding merits by the Italian Ministry of University and Research.
Expert reviewer for the TMLR journal.
Outstanding reviewer at NeurIPS '21, '22, '23, '24, ICLR '22, '24, ICML '22, '24.

Editorial Activities

Program Co-Chair of the European Workshop on Reinforcement Learning (EWRL 2022).

Action Editor of the TMLR journal 2023, 2024

Reviewer of

NeurIPS 2021*, 2022*, 2023*, 2024*

*Outstanding reviewer

ICML 2021, 2022*, 2023, 2024*

†Expert reviewer

ICLR 2021, 2022*, 2023, 2024*, 2025

AAAI 2021, 2022, 2023

AISTATS 2021, 2022

IJCAI 2020, 2021

TMLR 2022†, 2023†

Machine Learning, JAIR, Neural Networks, IEEE RA-L

Invited Talks

- 04/2024 **Unsupervised Reinforcement Learning**, VANDAL lab at Politecnico di Torino.
- 06/2023 **(Non)Convex Reinforcement Learning**, LAS group & AI Center at ETH Zurich.
- 02/2023 **(Non)Convex Reinforcement Learning**, Rising Stars in AI Symposium 2023 at KAUST.
- 09/2022 **Unsupervised Reinforcement Learning**, Robot Learning Lab at the Technion.
- 05/2022 **Unsupervised Reinforcement Learning**, Facebook AI Research - External Seminars.
- 11/2019 **Configurable Markov Decision Processes**, International Conference of the Italian Association for Artificial Intelligence (AlxIA).

Student Co-Supervision

- 2024 **Yonatan Ashlag**, M.Sc. in Electrical and Computing Engineering, Technion.
- 2023-2024 **Duilio Cirino**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2023 **Stefano Civelli**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2022-2023 **Vlad Marian Cimpeanu**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2021-2023 **Majid Molaei**, M.Sc. in Telecommunications Engineering, Politecnico di Milano.
- 2021-2022 **Juan Calderon**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2021-2022 **Pietro Maldini**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2020-2021 **Mattia Mancassola**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2020-2021 **Stefano Del Col**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2019-2023 **Riccardo De Santi**, M.Sc. in Computer Science, ETH Zurich.
- 2019-2020 **Lorenzo Pratissoli**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.
- 2018-2019 **Giovanni Pelosi**, M.Sc. in Computer Science and Engineering, Politecnico di Milano.

Publications

Journal

- [J1] **Mirco Mutti**, Riccardo De Santi, Piersilvio De Bartolomeis, and Marcello Restelli. *Convex Reinforcement Learning in Finite Trials*. JMLR, 2023.

Conference

- [18] Chinmaya Kausik, **Mirco Mutti**, Aldo Pacchiano, Ambuj Tewari. *A Theoretical Framework for Partially-Observed Reward-States in RLHF*. ICLR 2025.

- [17] Filippo Lazzati **Mirco Mutti**, and Alberto Maria Metelli. *How does Inverse RL Scale to Large State Spaces? A Provably Efficient Approach*. NeurIPS 2024.
- [16] Riccardo Zamboni, Duilio Cirino, Marcello Restelli, and **Mirco Mutti**. *The Limits of Pure Exploration in POMDPs: When the Observation Entropy is Enough*. RLC 2024.
- [15] Riccardo Zamboni, Duilio Cirino, Marcello Restelli, and **Mirco Mutti**. *How to Explore with Belief: State Entropy Maximization in POMDPs*. ICML 2024.
- [14] **Mirco Mutti** and Aviv Tamar. *Test-Time regret Minimization in Meta Reinforcement Learning*. ICML 2024.
- [13] Riccardo De Santi, Federico Arangath Joseph, Noah Liniger, **Mirco Mutti**, and Andreas Krause. *Geometric Active Exploration in Markov Decision Processes: the Benefit of Abstraction*. ICML 2024.
- [12] Filippo Lazzati, **Mirco Mutti**, and Alberto Maria Metelli. *Offline Inverse RL: New Solution Concepts and Provably Efficient Algorithms*. ICML 2024.
- [11] **Mirco Mutti**, Riccardo De Santi, Marcello Restelli, Alexander Marx, and Giorgia Ramponi. *Exploiting Causal Graph Priors with Posterior Sampling for Reinforcement Learning*. ICLR 2024.
- [10] Martino Bernasconi, Matteo Castiglioni, Alberto Marchesi, and **Mirco Mutti**. *Persuading Farsighted receivers in MDPs: the Power of Honesty*. NeurIPS 2023.
- [9] Alberto Maria Metelli, **Mirco Mutti**, and Marcello Restelli. *A Tale of Sampling and Estimation in Discounted Reinforcement Learning*. AISTATS 2023. Notable paper: 32/1689 (2%).
- [8] **Mirco Mutti**, Riccardo De Santi, Emanuele Rossi, Juan Felipe Calderon, Michael Bronstein, and Marcello Restelli. *Provably Efficient Causal Model-Based Reinforcement Learning for Systematic Generalization*. AAAI 2023.
- [7] **Mirco Mutti**, Riccardo De Santi, Piersilvio De Bartolomeis, and Marcello Restelli. *Challenging Common Assumptions in Convex Reinforcement Learning*. NeurIPS 2022.
- [6] **Mirco Mutti**, Riccardo De Santi, and Marcello Restelli. *The Importance of Non-Markovianity in Maximum State Entropy Exploration*. ICML 2022. Outstanding paper award: 10/5630 (0.17%).
- [5] **Mirco Mutti**, Stefano Del Col, and Marcello Restelli. *Reward-Free Policy Space Compression for Reinforcement Learning*. AISTATS 2022.
- [4] **Mirco Mutti**, Mattia Mancassola, and Marcello Restelli. *Unsupervised Reinforcement Learning in Multiple Environments*. AAAI 2022.
- [3] **Mirco Mutti**, Lorenzo Pratissoli, and Marcello Restelli. *Task-Agnostic Exploration via Policy Gradient of a Non-Parametric State Entropy Estimate*. AAAI 2021.
- [2] **Mirco Mutti** and Marcello Restelli. *An Intrinsically-Motivated Approach for Learning Highly Exploring and Fast Mixing Policies*. AAAI 2020.
- [1] Alberto Maria Metelli*, **Mirco Mutti***, and Marcello Restelli. *Configurable Markov Decision Processes*. ICML 2018. Long talk: 210/2473 (8.5%). *equal contribution

Preprint

- [P1] Filippo Lazzati, **Mirco Mutti**, Alberto Maria Metelli. *Reward Compatibility: A Framework for Inverse RL*. Arxiv 2025.

Selected workshop

- [W4] **Mirco Mutti**, Riccardo De Santi, Marcello Restelli, Alexander Marx, and Giorgia Ramponi. *Exploiting Causal Graph Priors with Posterior Sampling for Reinforcement Learning*. EWRL 2023.
- [W3] **Mirco Mutti**, Riccardo De Santi, Emanuele Rossi, Juan Felipe Calderon, Michael Bronstein, and Marcello Restelli. *Provably Efficient Causal Model-Based Reinforcement Learning for Systematic Generalization*. RLDM 2022.
- [W2] **Mirco Mutti** and Marcello Restelli. *An Approach for Learning Highly Exploring and Fast Mixing Policies via Intrinsic Motivation*. Workshop on Intrinsically Motivated Open-ended Learning (IMOL) 2019.
- [W1] Alberto Maria Metelli*, **Mirco Mutti***, and Marcello Restelli. *Configurable Markov Decision Processes*. EWRL 2018. *equal contribution