# Marco Mussi | Assistant Professor

marco.mussi@polimi.it • marcomussi.github.io • marcomussi95 • marcomussi

## Summary

I am Assistant Professor within the Dipartimento di Elettronica, Informazione e Bioingegneria, in the Artificial Intelligence and Robotic Laboratory of Politecnico di Milano. I received the Doctor of Philosophy in Information Technology with honors at Politecnico di Milano in June 2024, supervised by Professor Marcello Restelli. My main research topics revolve around artificial intelligence and machine learning, focusing on online learning and reinforcement learning both from theoretical and applicative perspectives. I have contributed to several industrial research projects funded by both private companies and public institutions. I am member of the ELLIS Society.

## Highlights

- Currently **Assistant Professor** with the Department of Electronics, Information and Bioengineering of **Politecnico di Milano**.
- **Doctor of Philosophy in Information Technology** cum Laude from Politecnico di Milano in June 2024.
- Author of **7 publications in peer-reviewed international journals** (all Q1 Scimago and all as main contributor). Author of **18 publications in peer-reviewed international conferences** (7 as main contributor), including ICML (13), NeurIPS (2), KDD (1), and AISTATS (1). According to Core Conference Rating: 16 publications in A* venues and 1 publication in A venues.
- **Two *Spotlights*** at ICML 2024 (A* Core, 335/9473 submissions – top 3.5%), **one *Oral*** at KDD 2022 (A* Core, 54/753 submissions – top 7%).
- Advisor of **1 Ph.D. student**, co-supervisor of **4 Ph.D. students** (all ongoing).
- Co-supervisor of **25 M.Sc. students theses** (23 concluded, 2 ongoing).
- **Lecturer** of the Ph.D. course on **Learning Theory** and of the B.Sc. course on **Computer Science** for Civil Engineering, both at Politecnico di Milano. **Teaching Assistant** of the M.Sc. course of **Machine Learning** at Politecnico di Milano. **Teaching Assistant** and **Laboratory Teaching Assistant** for several B.Sc. courses of **Computer Science**.
- **Co-organizer** of the 15th **European Workshop on Reinforcement Learning** (EWRL 2022). **Organizer** and **Program Chair** of the 2nd **Workshop on Aligning Reinforcement Learning Experimentalists and Theorists** (ARLET @ NeurIPS 2025) and of the **Workshop on AI for Safety-Critical Infrastructures** (AI-SCI @ ECML 2025). **Co-organizer** of the **Reinforcement Learning Summer School** (RLSS 2026).
- **Area Chair** of NeurIPS and KDD. **Program Committee** member of ICML, AAAI, ICLR, and AISTATS. **Reviewer** for Machine Learning (Springer), Transactions on Neural Networks and Learning Systems (IEEE), Robotics and Automation Letters (IEEE), and Engineering Applications of Artificial Intelligence (Elsevier).
- Participation in **5 industrial research** projects funded by private companies and public research institutions.
- Participation in the **Horizon Europe** competitive research project **AI4REALNET** (AI for Real-World Network Operation), and in the **FIS** competitive research project **HUmLrn** (Unified Learning from Diverse Human Feedback).
- **Winner** of the Dimitri N. Chorafas Award, recognized to the two **best Ph.D. Thesis** of the year in **Information Technology** at Politecnico di Milano ($\approx 100$ candidates).
- **Member** of the **ELLIS Society**.

## Education

**Ph.D. in Information Technology (Dottorato di Ricerca)** — **Milano**
*Politecnico di Milano – Dipartimento di Elettronica, Informazione e Bioingegneria* — *Nov 2020 – Jun 2024*
*Advisor:* Prof. Marcello Restelli
*Thesis:* Online Learning Methods for Pricing and Advertising
*Reviewers:* Alessandro Lazaric (Meta AI, Paris, France) and Tom Cesari (University of Ottawa, Ottawa, Canada)
*Final Mark:* Laude
*Date of Award:* June 27, 2024

**M.Sc. in Computer Science and Engineering (LM-32 Ingegneria Informatica)** — **Milano**
*Politecnico di Milano* — *Sep 2017 – Dec 2019*
*Main focus:* Artificial Intelligence and Machine Learning
*Thesis:* Improving Aerodynamic Load Estimation Algorithms for F1 Racing Cars
*Supervisor:* Prof. Marcello Restelli
*Industrial Partner:* Scuderia Ferrari F1
*Date of Award:* December 18, 2019

**B.Sc. in Engineering of Computing Systems (L-8 Ingegneria dell'Informazione)** — **Milano**
*Politecnico di Milano* — *Sep 2014 – Jul 2017*
*Date of Award:* July 25, 2017

**High School Diploma in Computer Science (Perito Informatico)** — **Crema**
*IIS Galileo Galilei* — *Sep 2008 – Jul 2014*

## Academic Experience

**Assistant Professor (Ricercatore RTD-A – L.240/2010, art.24, c.3, lett.A)** — **Milano**
*Politecnico di Milano* — *Jan 2026 – now*
*SSD:* IINF-05/A, *GSD:* 09/IINF-05 – Sistemi di Elaborazione delle Informazioni

**Postdoctoral Researcher (Assegno di Ricerca – L.240/2010, art.22)** — **Milano**
*Politecnico di Milano* — *Jun 2024 – Jan 2026*
*Title:* Development of machine learning algorithms for the analysis of extreme climate events
*SSD:* IINF-05/A – Sistemi di Elaborazione delle Informazioni
*Research Manager:* Prof. Marcello Restelli

**Research Assistant (Assegno di Ricerca – L.240/2010, art.22)** — **Milano**
*Politecnico di Milano* — *Jan 2020 – Oct 2020*
*Title:* Development of machine learning algorithms for diagnostics on lithium-ion batteries
*SSD:* ING-INF/05 – Sistemi di Elaborazione delle Informazioni
*Research Manager:* Prof. Marcello Restelli

## Publications

### Overview

- **7 publications in peer-reviewed international journals**, 7/7 as main contributor, 7/7 in Q1 journals according to Scimago, including Artificial Intelligence (AIJ) and Journal of Artificial Intelligence Research (JAIR).
- **18 publications in peer-reviewed international conferences**, 7 as main contributor, including ICML (13), NeurIPS (2), KDD (1), and AISTATS (1). According to Core Conference Rating: 16 publications in A* venues and 1 publication in A venues.
- **13 publications in peer-reviewed international workshops**.

### Publications' Pages

| | |
|---|---|
| **Google Scholar** | `scholar.google.com/citations?user=3gca-JUAAAAJ&hl=en` |
| **Scopus** | `scopus.com/authid/detail.uri?authorId=57407218600` |
| **DBLP** | `dblp.org/pid/321/0756` |

## List of Publications

The first author represent the main contributor. In the case of multiple main contributors (equal contribution), they are marked using the * symbol. For both journal and conference publications, the rating is the most recent available at the moment of publication.

| | Article Type | Ranking Portal | Link |
|---|---|---|---|
| Rankings' source: | Conference Proceedings | Core | portal.core.edu.au/conf-ranks |
| | Journals | Scimago | scimagojr.com |
| | | Core | portal.core.edu.au/jnl-ranks |

## Book Chapters

[B1] Marco Mussi. Multi-Armed Bandits Algorithms for Pricing and Advertising. Special Topics in Information Technology, Springer International Publishing, 2026.
Link: https://doi.org/10.1007/978-3-032-12359-6_3

## Refereed International Journal Articles

[J1] Marco Mussi, Alberto Maria Metelli, Marcello Restelli, Gianvito Losapio, Ricardo Jorge Bessa, Daniel Boos, Clark Borst, Giulia Leto, Alberto Castagna, Ricardo Chavarriaga, Duarte Dias, Adrian Egli, Andrina Eisenegger, Yassine El Manyari, Anton Fuxjäger, Joaquim Geraldes, Samira Hamouche, Mohamed Hassouna, Bruno Lemetayer, Milad Leyli-Abadi, Roman Liessner, Jonas Lundberg, Antoine Marot, Maroua Meddeb, Viola Schiaffonati, Manuel Schneider, Thilo Stadelmann, Julia Usher, Herke van Hoof, Jan Viebahn, Toni Waefler and Giacomo Zanotti. Human-AI Interaction in Safety-Critical Network Infrastructures. iScience, volume 28, 113400, Cell Press, 2025.
Scimago 2024: Q1 (Multidisciplinary)
Link: https://doi.org/10.1016/j.isci.2025.113400

[J2] Marco Mussi*, Simone Drago*, Marcello Restelli and Alberto Maria Metelli. Factored-Reward Bandits with Intermediate Observations: Regret Minimization and Best Arm Identification. Artificial Intelligence, volume 347, 104362, Elsevier, 2025.
Scimago 2024: Q1 (Computer Science) – Core 2020: A*
Link: https://doi.org/10.1016/j.artint.2025.104362

[J3] Marco Mussi and Alberto Maria Metelli. Generalizing the Regret: an Analysis of Lower and Upper Bounds. Journal of Artificial Intelligence Research, volume 82, pages 1773–1806, 2025.
Scimago 2023: Q1 (Computer Science) – Core 2020: A
Link: https://doi.org/10.1613/jair.1.17614

[J4] Marco Mussi, Luigi Pellegrino, Oscar Francesco Pindaro, Marcello Restelli and Francesco Trovó. A Reinforcement Learning Controller Optimizing Costs and Battery State of Health in Smart Grids. Journal of Energy Storage, volume 82, pages 110572–110580, Elsevier, 2024.
Scimago 2022: Q1 (Engineering)
Link: https://doi.org/10.1016/j.est.2024.110572

[J5] Marco Mussi, Davide Lombarda, Alberto Maria Metelli, Francesco Trovó and Marcello Restelli. ARLO: A Framework for Automated Reinforcement Learning. Expert Systems with Applications, volume 224, pages 119883–119894, Elsevier, 2023.
Scimago 2022: Q1 (Computer Science and Engineering) – Core 2020: B
Link: https://doi.org/10.1016/j.eswa.2023.119883

[J6] Marco Mussi, Luigi Pellegrino, Marcello Restelli and Francesco Trovó. An Online State of Health Estimation Method for Lithium-Ion Batteries based on Time Partitioning and Data-Driven Model Identification. Journal of Energy Storage, volume 55, pages 105467–105474, Elsevier, 2022.
Scimago 2021: Q1 (Engineering)
Link: https://doi.org/10.1016/j.est.2022.105467

[J7] Marco Mussi, Luigi Pellegrino, Marcello Restelli and Francesco Trovó. A voltage dynamic-based state of charge estimation method for batteries storage systems. Journal of Energy Storage, volume 44, pages 103309–103318, Elsevier, 2021.
Scimago 2020: Q1 (Engineering)
Link: https://doi.org/10.1016/j.est.2021.103309

## Refereed International Conferences Papers

[C1] Simone Drago, Marco Mussi and Alberto Maria Metelli. Online Compatible Reward Identification from Preference Feedback. Proceedings of the 43rd International Conference on Machine Learning (ICML), 2026. To Appear.

Acceptance rate 26.6% (6352/23918) – Core 2026: A* – GGS 2021: A++
Link: To Appear.

[C2] Federico Corso, Marco Mussi and Alberto Maria Metelli. Fast Mixing Steady-State Control in Markov Decision Processes. Proceedings of the 43rd International Conference on Machine Learning (ICML), 2026. To Appear.
Acceptance rate 26.6% (6352/23918) – Core 2026: A* – GGS 2021: A++
Link: To Appear.

[C3] Cristiano Migali, Gianmarco Genalti, Alberto Maria Metelli and Marco Mussi. Learning to Rank from Incomplete Rankings. Proceedings of the 43rd International Conference on Machine Learning (ICML), 2026. To Appear.
Acceptance rate 26.6% (6352/23918) – Core 2026: A* – GGS 2021: A++
Link: To Appear.

[C4] Alessandro Montenegro, Federico Mansutti, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Reusing Trajectories in Policy Gradients Enables Fast Convergence. Proceedings of the 43rd International Conference on Machine Learning (ICML), 2026. To Appear.
Acceptance rate 26.6% (6352/23918) – Core 2026: A* – GGS 2021: A++
Link: To Appear.

[C5] Cristiano Migali, Marco Mussi, Gianmarco Genalti and Alberto Maria Metelli. Tightening Regret Lower and Upper Bounds in Restless Rising Bandits. Advances in Neural Information Processing Systems (NeurIPS), volume 38, pages 18437–18478, 2025.
Acceptance rate 24.5% (5290/21575) – Core 2023: A* – GGS 2021: A++
Link: `https://openreview.net/forum?id=9tCJHfF6M4`

[C6] Simone Drago*, Marco Mussi* and Alberto Maria Metelli. Sleeping Reinforcement Learning. Proceedings of the 42nd International Conference on Machine Learning (ICML), volume 267, pages 14439–14498, PMLR, 2025.
Acceptance rate 26.9% (3260/12107) – Core 2023: A* – GGS 2021: A++
Link: `https://proceedings.mlr.press/v267/drago25a`

[C7] Simone Drago, Marco Mussi and Alberto Maria Metelli. Towards Theoretical Understanding of Sequential Decision Making with Preference Feedback. Proceedings of the 42nd International Conference on Machine Learning (ICML), volume 267, pages 14499–14514, PMLR, 2025.
Acceptance rate 26.9% (3260/12107) – Core 2023: A* – GGS 2021: A++
Link: `https://proceedings.mlr.press/v267/drago25b`

[C8] Alessandro Montenegro, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Convergence Analysis of Policy Gradient Methods with Dynamic Stochasticity. Proceedings of the 42nd International Conference on Machine Learning (ICML), volume 267, pages 44652–44698, PMLR, 2025.
Acceptance rate 26.9% (3260/12107) – Core 2023: A* – GGS 2021: A++
Link: `https://proceedings.mlr.press/v267/montenegro25a`

[C9] Simone Drago, Marco Mussi and Alberto Maria Metelli. Position: Constants are Critical in Regret Bounds for Reinforcement Learning. Proceedings of the 42nd International Conference on Machine Learning (ICML), Position Papers Track, volume 267, pages 81233–81275, PMLR, 2025.
Acceptance rate 19.7% (71/361) – Core 2023: A* – GGS 2021: A++
Link: `https://proceedings.mlr.press/v267/drago25c`

[C10] Alessandro Montenegro, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Last-Iterate Global Convergence of Policy Gradients for Constrained Reinforcement Learning. Advances in Neural Information Processing Systems (NeurIPS), volume 37, pages 126363–126416, 2024.
Acceptance rate 25.8% (4039/15671) – Core 2023: A* – GGS 2021: A++
Link: `https://openreview.net/forum?id=2vywag2lVC`

[C11] Marco Mussi*, Simone Drago*, Marcello Restelli and Alberto Maria Metelli. Factored-Reward Bandits with Intermediate Observations. Proceedings of the 41st International Conference on Machine Learning (ICML), volume 235, pages 36911–36952, PMLR, 2024.
Acceptance rate 27.5% (2609/9473) – Core 2023: A* – GGS 2021: A++
Link: `https://proceedings.mlr.press/v235/mussi24a`

[C12] Marco Mussi, Alessandro Montenegro, Francesco Trovò, Marcello Restelli and Alberto Maria Metelli. Best Arm Identification for Stochastic Rising Bandits. Proceedings of the 41st International Conference on

Machine Learning (ICML), volume 235, pages 36953–36989, PMLR, 2024.
Spotlight: top 3.5% (335/9473) – Core 2023: A* – GGS 2021: A++
Link: https://proceedings.mlr.press/v235/mussi24b

[C13] Alessandro Montenegro, Marco Mussi, Alberto Maria Metelli and Matteo Papini. Learning Optimal Deterministic Policies with Stochastic Policy Gradients. Proceedings of the 41st International Conference on Machine Learning (ICML), volume 235, pages 36160–36211, PMLR, 2024.
Spotlight: top 3.5% (335/9473) – Core 2023: A* – GGS 2021: A++
Link: https://proceedings.mlr.press/v235/montenegro24a

[C14] Gianmarco Genalti, Marco Mussi, Nicola Gatti, Marcello Restelli, Matteo Castiglioni and Alberto Maria Metelli. Graph-Triggered Rising Bandits. Proceedings of the 41st International Conference on Machine Learning (ICML), volume 235, pages 15351–15380, PMLR, 2024.
Acceptance rate 27.5% (2609/9473) – Core 2023: A* – GGS 2021: A++
Link: https://proceedings.mlr.press/v235/genalti24a

[C15] Francesco Bacchiocchi*, Gianmarco Genalti*, Davide Maran*, Marco Mussi*, Marcello Restelli, Nicola Gatti and Alberto Maria Metelli. Autoregressive Bandits. Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS), volume 238, pages 937–945, PMLR, 2024.
Acceptance rate 27.6% (546/1980) – Core 2023: A – GGS 2021: A+
Link: https://proceedings.mlr.press/v238/bacchiocchi24a

[C16] Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Dynamical Linear Bandits. Proceedings of the 40th International Conference on Machine Learning (ICML), volume 202, pages 25563–25587, PMLR, 2023.
Acceptance rate 27.9% (1827/6538) – Core 2023: A* – GGS 2021: A++
Link: https://proceedings.mlr.press/v202/mussi23a

[C17] Marco Mussi*, Gianmarco Genalti*, Alessandro Nuara, Francesco Trovó, Marcello Restelli and Nicola Gatti. Dynamic Pricing with Volume Discounts in Online Settings. Proceedings of the 35th Conference on Innovative Applications of Artificial Intelligence (IAAI), volume 37, pages 15560–15568, AAAI, 2023. Innovative Application of AI Award.
Core 2021: B – GGS 2021: B-
Link: https://doi.org/10.1609/aaai.v37i13.26845

[C18] Marco Mussi, Gianmarco Genalti, Francesco Trovó, Alessandro Nuara, Nicola Gatti and Marcello Restelli. Pricing the Long Tail by Explainable Product Aggregation and Monotonic Bandits. Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pages 3623–3633, ACM, 2022.
Oral Presentation: top 7% (54/753) – Core 2021: A* – GGS 2021: A++
Link: https://doi.org/10.1145/3534678.3539142

## Refereed International Workshop Papers

[W1] Federico Corso, Marco Mussi and Alberto Maria Metelli. Trading-off Reward Maximization and Stability in Sequential Decision Making. European Workshop on Reinforcement Learning (EWRL), 2025.
Link: https://openreview.net/forum?id=5FSTwSgQQz

[W2] Simone Drago, Marco Mussi and Alberto Maria Metelli. A Theoretical Perspective on Sequential Decision Making with Preference Feedback. European Workshop on Reinforcement Learning (EWRL), 2025.
Link: https://openreview.net/forum?id=eD4GO3mSvL

[W3] Alberto Maria Metelli, Simone Drago and Marco Mussi. A Novel Self-Normalized Bernstein-Like Dimension-Free Inequality and Regret Bounds for Generalized Kernelized Bandits. European Workshop on Reinforcement Learning (EWRL), 2025.
Link: https://openreview.net/forum?id=kHglhH6Sd9

[W4] Davide Salaorni, Vincenzo De Paola, Samuele Delpero, Giovanni Dispoto, Paolo Bonetti, Alessio Russo, Giuseppe Calcagno, Francesco Trovó, Matteo Papini, Alberto Maria Metelli, Marco Mussi and Marcello Restelli. Gym4ReaL: A Benchmark Suite for Evaluating Reinforcement Learning in Realistic Domains. European Workshop on Reinforcement Learning (EWRL), 2025.
Link: https://openreview.net/forum?id=Aomqn4nAhK

[W5] Carlo Fabrizio*, Gianvito Losapio*, Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Power Grid Control with Graph-Based Distributed Reinforcement Learning. Workshop on Machine Learning for Sustainable Power Systems at the European Conference on Machine Learning (ECML), 2025.

Oral Presentation
Link: https://marcomussi.github.io/papers/ecmlw2025/paper.pdf

[W6] Gianvito Losapio, Davide Beretta, Marco Mussi, Alberto Maria Metelli and Marcello Restelli. State and Action Factorization in Power Grids. Workshop on Machine Learning for Sustainable Power Systems at the European Conference on Machine Learning (ECML), 2024.
Oral Presentation
Link: https://marcomussi.github.io/papers/ecmlw2024/paper.pdf

[W7] Simone Drago and Marco Mussi. Open Problem: Tight Bounds for Bernoulli Rewards in Kernelized Multi-Armed Bandits. Workshop on Aligning Reinforcement Learning Experimentalists and Theorists at the International Conference on Machine Learning (ICML), 2024.
Link: https://marcomussi.github.io/papers/icmlarlet2024/paper.pdf

[W8] Simone Drago, Marco Mussi, Marcello Restelli and Alberto Maria Metelli. Intermediate Observations in Factored-Reward Bandits. Adaptive and Learning Agents Workshop at the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2024.
Link: https://openreview.net/forum?id=CIyAiUCfrj

[W9] Francesco Bacchiocchi*, Gianmarco Genalti*, Davide Maran*, Marco Mussi*, Marcello Restelli, Nicola Gatti and Alberto Maria Metelli. Online Learning in Autoregressive Dynamics. European Workshop on Reinforcement Learning (EWRL), 2023.
Link: https://openreview.net/forum?id=YrHpQWpwsy

[W10] Alessandro Montenegro, Marco Mussi, Francesco Trovò, Marcello Restelli and Alberto Maria Metelli. Stochastic Rising Bandits: A Best Arm Identification Approach. European Workshop on Reinforcement Learning (EWRL), 2023.
Link: https://openreview.net/forum?id=Ctq0d9LEuT

[W11] 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 the International Conference on Machine Learning (ICML), 2023.
Link: https://openreview.net/forum?id=k6aftfkuad

[W12] Gianmarco Genalti, Marco Mussi, Alessandro Nuara and Nicola Gatti. Dynamic Pricing with Online Data Aggregation and Learning. European Workshop on Reinforcement Learning (EWRL), 2022.
Oral Presentation: top 10.5% (10/96)
Link: https://marcomussi.github.io/papers/ewrl2022/paper.pdf

[W13] Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Dynamical Linear Bandits for Long-Lasting Vanishing Rewards. Complex Feedback in Online Learning Workshop at the International Conference on Machine Learning (ICML), 2022.
Link: https://marcomussi.github.io/papers/icml2022/paper.pdf

## Technical Reports

[R1] Marco Mussi, Gianvito Losapio, Alberto Maria Metelli, Marcello Restelli, Ricardo Bessa, Antoine Marot, Daniel Boos, Clark Borst, Alberto Castagna, Duarte Dias, Adrian Egli, Andrina Eisenegger, Yassine El Manyari, Anton Fuxjäger, Samira Hamouche, Mohamed Hassouna, Bruno Lemetayer, Roman Liessner, Jonas Lundberg, Manuel Schneider, Irene Sturm, Julia Usher, Herke Van Hoof, Jan Viebahn, Toni Wäfler. Position paper on AI for the operation of critical energy and mobility network infrastructures. AI4REALNET, 2024.
Link: https://ai4realnet.eu/deliverables

## Theses

[T1] Marco Mussi. Online Learning Methods for Pricing and Advertising. Doctoral Thesis. Politecnico di Milano. 2024.
Link: https://hdl.handle.net/10589/221772

[T2] Marco Mussi. Improving Aerodynamic Load Estimation Algorithms for F1 Racing Cars. Master's Thesis. Politecnico di Milano. 2019.
Link: https://hdl.handle.net/10589/152241

## Preprints

[P1] Marco Mussi and Marcello Restelli. Online Dynamic Pricing of Complementary Products. arXiv preprint, arXiv:2511.22291. 2025.
Link: `https://doi.org/10.48550/arXiv.2511.22291`

[P2] Alberto Maria Metelli, Simone Drago and Marco Mussi. Generalized Kernelized Bandits: A Novel Self-Normalized Bernstein-Like Dimension-Free Inequality and Regret Bounds. arXiv preprint, arXiv:2508.01681. 2025.
Link: `https://doi.org/10.48550/arXiv.2508.01681`

[P3] Davide Salaorni, Vincenzo De Paola, Samuele Delpero, Giovanni Dispoto, Paolo Bonetti, Alessio Russo, Giuseppe Calcagno, Francesco Trovó, Matteo Papini, Alberto Maria Metelli, Marco Mussi and Marcello Restelli. Gym4ReaL: A Suite for Benchmarking Real-World Reinforcement Learning. arXiv preprint, arXiv:2507.00257. 2025.
Link: `https://doi.org/10.48550/arXiv.2507.00257`

[P4] Alessandro Montengro, Leonardo Cesani, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Learning Deterministic Policies with Policy Gradients in Constrained Markov Decision Processes. arXiv preprint, arXiv:2506.05953. 2025.
Link: `https://doi.org/10.48550/arXiv.2506.05953`

[P5] Simone Drago, Marco Mussi and Alberto Maria Metelli. A refined Analysis of UCBVI. arXiv preprint, arXiv:2502.17370. 2025.
Link: `https://doi.org/10.48550/arXiv.2502.17370`

[P6] Gianmarco Genalti, Marco Mussi, Nicola Gatti, Marcello Restelli, Matteo Castiglioni and Alberto Maria Metelli. Bridging Rested and Restless Bandits with Graph-Triggering: Rising and Rotting. arXiv preprint, arXiv:2409.05980. 2024.
Link: `https://doi.org/10.48550/arXiv.2409.05980`

[P7] Marco Mussi, Simone Drago and Alberto Maria Metelli. Open Problem: Tight Bounds for Kernelized Multi-Armed Bandits with Bernoulli Rewards. arXiv preprint, arXiv:2407.06321. 2024.
Link: `https://doi.org/10.48550/arXiv.2407.06321`

## Under Review

[U1] Gianmarco Genalti, Marco Mussi, Nicola Gatti, Marcello Restelli, Matteo Castiglioni and Alberto Maria Metelli. Bridging Rested and Restless Bandits with Graph-Triggering: Rising and Rotting. 2024. Under Review at the Journal of Machine Learning Research (JMLR).

[U2] Gianmarco Tedeschi, Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Trading-off Statistical and Computational Efficiency via $W$-step MDPs: A Policy Gradient Approach. 2025. Under Review at the Machine Learning Journal (MLJ).

[U3] Cristiano Migali, Marco Mussi, Gianmarco Genalti and Alberto Maria Metelli. Near-Optimal Regret Lower and Upper Bounds for Restless Rising Bandits. 2025. Under Review at the Journal of Machine Learning Research (JMLR).

[U4] Alberto Maria Metelli, Simone Drago and Marco Mussi. Generalized Kernelized Bandits: A Novel Self-Normalized Bernstein-Like Dimension-Free Inequality and Regret Bounds. 2025. Under Review at the Journal of Machine Learning Research (JMLR).

[U5] Federico Corso, Riccardo Zamboni, Marco Mussi, Marcello Restelli and Alberto Maria Metelli. No-regret Learning with Revealed Transitions in Adversarial Markov Decision Processes. 2024. Under Review at the Artificial Intelligence Journal (AIJ).

[U6] Alessandro Montenegro, Leonardo Cesani, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Learning Deterministic Policies with Policy Gradients in Constrained Markov Decision Processes. 2025. Under Review at the Artificial Intelligence Journal (AIJ).

[U7] Marco Mussi and Marcello Restelli. Online Dynamic Pricing of Complementary Products. 2025. Under Review at the Artificial Intelligence Journal (AIJ).

[U8] Simone Drago, Marco Mussi and Alberto Maria Metelli. On the Computational Limits of Sequential Decision Making with Trajectory-Preference Feedback. 2025. Under Review at Algorithmica.

[U9] Davide Salaorni, Vincenzo De Paola, Samuele Delpero, Giovanni Dispoto, Paolo Bonetti, Alessio Russo, Giuseppe Calcagno, Francesco Trovó, Matteo Papini, Alberto Maria Metelli, Marco Mussi and Marcello Restelli. Gym4ReaL: Towards Real-World Reference Environments for Reinforcement Learning. 2025. Under Review at the ACM International Conference on Knowledge Discovery and Data Mining (KDD 2026).

[U10] Simone Drago, Marco Mussi and Alberto Maria Metelli. Generalizing Preference-based Reinforcement Learning: A Rationality Model for Incomparability. 2026. Under Review at the International Conference on Neural Information Processing Systems (NeurIPS 2026).

[U11] Marco Mussi, Marcello Restelli and Alberto Maria Metelli. Computationally Efficient and Stable Gaussian Processes with Controllable Approximation Error for Supervised and Online Learning. 2026. Under Review at the International Conference on Neural Information Processing Systems (NeurIPS 2026).

[U12] Diego Alovisetti, Marco Mussi and Alberto Maria Metelli. Non-parametric Ranking in Log-concave Random Utility Models. 2026. Under Review at the International Conference on Neural Information Processing Systems (NeurIPS 2026).

[U13] Cristiano Migali, Gianmarco Genalti, Alberto Maria Metelli and Marco Mussi. Worst-Case Regret Bounds for Combinatorial Bandits with Ranking Feedback. 2026. Under Review at the International Conference on Neural Information Processing Systems (NeurIPS 2026).

[U14] Alessandro Montenegro, Riccardo Venturelli, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Reusing Past Samples in Proximal Policy Optimization: When and How Does It Help?. 2026. Under Review at the International Conference on Neural Information Processing Systems (NeurIPS 2026).

### In Preparation

[N1] Simone Drago, Marco Mussi and Alberto Maria Metelli. Optimal Algorithm for Multi-Objective Reinforcement Learning. 2026. Expected Submission to the Annual AAAI Conference on Artificial Intelligence (AAAI 2027).

[N2] Federico Corso, Marco Mussi and Alberto Maria Metelli. On Stability of Reinforcement Learning Policies in Continuous State and Action Spaces. 2026. Expected Submission to the Annual AAAI Conference on Artificial Intelligence (AAAI 2027).

[N3] Alberto Maria Metelli, Gianmarco Tedeschi, Alessandro Montenegro and Marco Mussi. Policy Gradients with Bias. 2026. Expected Submission to the International Conference on Artificial Intelligence and Statistics (AISTATS 2027).

[N4] Marco Mussi, Marcello Restelli and Alberto Maria Metelli. Behavioral Cloning from Human Expert Feedback in Utility-based Markov Decision Processes. 2026. Expected Submission to the International Conference on Artificial Intelligence and Statistics (AISTATS 2027).

## Competitive Research Projects

**HUmLrn (Unified Learning from Diverse Human Feedback)** **Politecnico di Milano**
*Funding Institution: Italian Ministry of University and Research (FIS2 Founding)* *Jan 2026 – now*
*Budget:* 1,324,909.96€
*Duration:* 36 months
*Principal Investigator:* Prof. Alberto Maria Metelli
*Description:* The HUmLrn project investigates how AI agents can effectively integrate diverse forms of human feedback to improve learning beyond traditional demonstration-based methods. Inspired by human learning, it proposes a unified framework that accommodates multiple feedback types, such as demonstrations and corrections, within a single methodology. The project develops theoretical foundations and novel learning algorithms addressing the statistical challenges of this setting.

**AI4REALNET (AI for REAL-world NETwork operation)** **Politecnico di Milano**
*Funding Institution: European Union (Horizon Europe)* *Oct 2023 – now*
*Budget:* 3,999,976.25€ (456,250€ to Politecnico di Milano)
*Duration:* 42 months
*Principal Investigator:* Prof. Marcello Restelli
*Roles:*
- Research Scientist in the AI Fundamental Research Work Package (WP2).

- Task Leader for Task 2.2.
- Responsible for Deliverable 2.1 (Position Paper on AI for Safety Critical Infrastructures).

*Description:* AI4REALNET focuses on AI solutions for critical systems like electricity, railway, and air traffic management, enhancing human operations through simulation-based networks. Its goals are: $(i)$ developing trustworthy AI-driven decision-making with augmented cognition and hybrid human-AI learning, ensuring resilience, safety, and security, and $(ii)$ advancing and validating AI algorithms using open-source digital environments that emulate real-world system operations.
*Partners:* Politecnico di Milano, INESC TEC, SBB Swiss Federal Railways, Delft University of Technology, EnliteAI, Zurich University of Applied Sciences, University of Applied Sciences and Arts Northwestern Switzerland, NAV Portugal, Fraunhofer IEE, University of Kassel, Reseau de Transport d'Electricite, IRT SystemX, Deutsche Bahn, Linkoping University, Flatland Association, University of Amsterdam, TenneT.

## Industrial Projects

**ML cube (Politecnico di Milano's Spin-off)** **Milano**
*Life-Cycle-Management and Optimization of Machine Learning Algorithms in Real-time Biddings* *2021 – 2023*
*Role:* Research Scientist
*Principal Investigators:* Nicola Gatti, Marcello Restelli and Francesco Trovó
*Duration:* 24 Months
*Value:* 84,000€ + VAT

**Ricerca Sistema Energetico (RSE)** **Milano**
*Development of Machine Learning Algorithms for Diagnostics on Lithium-ion Batteries* *2020 – 2021*
*Role:* Research Scientist
*Principal Investigators:* Marcello Restelli and Francesco Trovó
*Duration:* 12 Months
*Value:* 35,000€ + VAT

**PaxMile (Politecnico di Milano's Spin-off)** **Milano**
*Optimization and Artificial Intelligence Algorithms for the PaxMile Last-mile Delivery System* *2020 – 2021*
*Role:* Research Scientist
*Principal Investigators:* Edoardo Amaldi and Marcello Restelli
*Duration:* 8 Months
*Value:* 50,000€ + VAT

**Ricerca Sistema Energetico (RSE)** **Milano**
*Development of Machine Learning Algorithms for Diagnostics on Lithium-ion Batteries* *2019 – 2020*
*Role:* Research Scientist
*Principal Investigator:* Marcello Restelli
*Duration:* 12 Months
*Value:* 35,000€ + VAT

**Scuderia Ferrari** **Milano**
*Improving Machine Learning Techniques for Aerodynamics Performance Optimization* *2019 – 2020*
*Role:* Student Researcher
*Principal Investigator:* Marcello Restelli
*Duration:* 12 Months
*Value:* 70,000€ + VAT

## Technology Transfer

### Products Release

**ADcube's Marketing Mix Model** **Milano**
*In collaboration with Politecnico di Milano's spin-off ML cube* *2021 – 2023*
*Focus:* Development and release of AD cube's Marketing Mix Model, a product for advertising optimization in online campaigns considering cross-channel interactions
*Role:* Development of all the Artificial Intelligence algorithms inside the product
*Funding:* Winner of the ELISE's 2nd Open Call Grants

**Dynamic Pricing for E-commerce** **Milano**
*In collaboration with the e-commerce website Euroffice* *2021 – 2022*
*Focus:* Development and release of a dynamic pricing solution for an e-commerce website managing over 20000 products
*Role:* Development of all the Artificial Intelligence algorithms inside the product

### Open Source Software Release

**Gym4ReaL: A Benchmarking Suite for Real-Life Reinforcement Learning** **Milano**
*Open Source Library to test Reinforcement Learning Algorithms* *2025 – now*
*Availability:* Available on GitHub

**ARLO: Automated Reinforcement Learning Optimizer** **Milano**
*Open Source Library for Automated Reinforcement Learning* *2021 – 2022*
*Availability:* Available on GitHub

## Teaching

### Lecturer

**Lecturer of Learning Theory** **15 Hours**
*Learning Theory – Politecnico di Milano* *Mar 2026 – Apr 2026*
Foundational aspects of machine learning and statistical learning theory.
Ph.D. Level – Information Technology.
A.Y. 2025-26. ≈50 students. 5 CFU.
Co-lecturer: Prof. Alberto Maria Metelli.
Course delivered in English.

**Lecturer of Computer Science** **36 Hours**
*Computer Science – Politecnico di Milano* *Feb 2026 – Jun 2026*
Foundational aspects of programming. Focus on `C` and `Fortran`.
B.Sc. in Civil Engineering – Milano Leonardo Campus.
A.Y. 2025-26. II semester. ≈190 students. 6 CFU.
Course delivered in English.
Students' Evaluation: Pending.

### Teaching Assistant

**Teaching Assistant of Machine Learning** **20 Hours**
*Machine Learning – Politecnico di Milano* *Feb 2026 – Jun 2026*
Exercise sessions mainly supervised and reinforcement learning.
M.Sc. in Computer Science and Engineering – Milano Leonardo Campus.
A.Y. 2025-26. II semester. ≈480 students. 5 CFU.
Lecturer: Prof. Marcello Restelli.
Course delivered in English.
Students' Evaluation: Pending.

**Teaching Assistant of Computer Science** **20 Hours**
*Foundations of Computer Science (Fondamenti di Informatica) – Politecnico di Milano* *Sep 2025 – Dec 2025*
Exercise sessions mainly on the `C` programming language.
B.Sc. in Engineering of Computing Systems – Milano Leonardo Campus.
A.Y. 2025-26. I semester. ≈270 students. 10 CFU.
Lecturer: Prof. Alberto Maria Metelli.
Course delivered in Italian.
Students' Evaluation: Pending.

**Teaching Assistant of Computer Science** **24 Hours**
*Computer Science – Politecnico di Milano* *Feb 2025 – Jun 2025*
Exercise sessions mainly on `C` and `Fortran` programming languages.
B.Sc. in Civil Engineering – Milano Leonardo Campus.
A.Y. 2024-25. II semester. ≈140 students. 6 CFU.
Lecturer: Prof. Alberto Maria Metelli.
Course delivered in English.
Students' Evaluation: 3.4/4 (University Average: 3.2/4).

**Teaching Assistant of Computer Science** **24 Hours**
*Computer Science – Politecnico di Milano* *Feb 2024 – Jun 2024*
Exercise sessions mainly on `C` and `Fortran` programming languages.
B.Sc. in Civil Engineering – Milano Leonardo Campus.
A.Y. 2023-24. II semester. ≈110 students. 6 CFU.
Lecturer: Prof. Alberto Maria Metelli.
Course delivered in English.
Students' Evaluation: 3.4/4 (University Average: 3.1/4).

### Laboratory Teaching Assistant

**Laboratory Teaching Assistant of Computer Science** **15 Hours**
*Computer Science (Informatica B) – Politecnico di Milano* *Sep 2024 – Dec 2024*
Laboratory sessions on `C` and `MATLAB` programming languages.
B.Sc. in Mechanical Engineering – Milano Bovisa Campus.
A.Y. 2024-25. I semester. ≈280 students. 7 CFU.
Lecturer: Prof. Francesco Trovó.
Course delivered in Italian.
Students' Evaluation: 3.1/4 (University Average: 3.2/4).

### Tutor

**Academic Tutor of Machine Learning** **30 Hours**
*Master in Artificial Intelligence and Machine Learning – Cefriel* *Sep 2022 – Jul 2023*
Supervision of a team in the application of Reinforcement Learning algorithms to real-world control problems.
Tutoring delivered in Italian.

## Member of Scientific Societies

**European Laboratory for Learning and Intelligent Systems (ELLIS)** **European Union**
*Member* *Jun 2025 – now*
*Focus:* ELLIS is a pan-European AI network founded in 2018. It promotes top-level research in machine learning to drive innovation and ensure Europe's leadership in AI.
Part also of the Ellis Unit Milan.

**Associazione Italiana per l'Intelligenza Artificiale (AIxIA)** **Italy**
*Member* *2025*
*Focus:* AIxIA has been promoting the study and research of Artificial Intelligence and coordinating activities in the field in Italy since 1988.

## Organization of International Events

**Reinforcement Learning Summer School (RLSS)** **Milano**
*Co-organizer*
Held in the Leonardo Campus of Politecnico di Milano from 3 to 12 June 2026.

**Aligning Reinforcement Learning Experimentalists and Theorists (ARLET) Workshop** **San Diego**
*Organizer and Program Chair*
Held within the Conference on Neural Information Processing Systems (NeurIPS 2025) on December 6, 2025.
$2^{nd}$ edition. Attendance: ≈150. Number of submitted papers: ≈140 (accepted: ≈100).
NeurIPS 2025 workshop proposals' acceptance rate: 19.4% (55/283).

**AI for Safety-Critical Infrastructures (AI-SCI) Workshop** **Porto**
*Co-organizer and Program Chair*
Held within the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD 2025) on September 15, 2025.

**European Workshop on Reinforcement Learning (EWRL)** **Milano**
*Co-organizer and Communication Chair*
Held in the Bovisa Campus of Politecnico di Milano from 19 to 21 September 2022.
$15^{th}$ edition. Attendance: ≈150. Number of submitted papers: ≈100 (accepted: ≈70).

## Collaborations

**AI4REALNET Consortium** **European Union**
*Various Research Collaborations* *2023 – now*
Collaborations in the context of the Horizon Europe project AI4REALNET with research centers and universities (Delft University of Technology, INESC TEC, Zurich University of Applied Sciences, University of Applied Sciences and Arts Northwestern Switzerland, Fraunhofer IEE, University of Kassel, IRT SystemX, Linköping University, University of Amsterdam) and companies (SBB Swiss Federal Railways, EnliteAI, NAV Portugal, Reseau de Transport d'Electricite, Deutsche Bahn, Flatland Association, TenneT). Leading role (first author) in the writing of a consortium paper (accepted).

**Politecnico di Milano** **Milano**
*Various Research Collaborations* *2020 – now*
Research collaborations with Prof. Marcello Restelli, Prof. Nicola Gatti, Prof. Francesco Trovò, Prof. Alberto Maria Metelli, Prof. Matteo Papini, Prof. Matteo Castiglioni, several postdoctoral researchers and PhD students.

# Invited Talks and Seminars

## Invited Talks

**This is IT! – Best Theses Talks (Ph.D. Students in Information Technology)** **Politecnico di Milano**
*Online Learning Methods for Pricing and Advertising* *19 November 2025*

## Seminars

**AI4REALNET Dissemination Seminar (Online) – Invited** **AI4REALNET Consortium**
*Distributed and Hierarchical Reinforcement Learning* *24 April 2024*

**DEIB Seminar** **Politecnico di Milano**
*An introduction to Reinforcement Learning in Real World* *3 September 2021*

**RSE Academy Seminar – Invited** **Ricerca Sistema Energetico**
*A Data-Driven Method for State of Charge Estimation* *23 October 2020*

# Awards and Recognitions

## International Conference Papers' Awards

**Spotlight Paper at the 41st International Conference on Machine Learning** **ICML 2024**
*Paper: "Learning Optimal Deterministic Policies with Stochastic Policy Gradients"* *July 2024*
Spotlight Paper: top 3.5% over 9473 submissions at ICML 2024 (Core 2023: A* – GGS 2021: A++).
Link: `https://icml.cc/virtual/2024/poster/34781`

**Spotlight Paper at the 41st International Conference on Machine Learning** **ICML 2024**
*Paper: "Best Arm Identification for Stochastic Rising Bandits"* *July 2024*
Spotlight Paper: top 3.5% over 9473 submissions at ICML 2024 (Core 2023: A* – GGS 2021: A++).
Link: `https://icml.cc/virtual/2024/poster/33841`

**Oral Presentation at the 28th Conference on Knowledge Discovery and Data Mining** **KDD 2022**
*Paper: "Pricing the Long Tail by Explainable Product Aggregation and Monotonic Bandits"* *August 2022*
Oral Presentation: top 7% over 753 submissions at KDD 2022 (Core 2021: A* – GGS 2021: A++).
Link: `https://kdd.org/kdd2022/paperADS.html`

## Other Papers' Awards

**Publications in Top Journals Award** **DEIB — Politecnico di Milano**
*Awarded to the Top-10 Papers from PoliMi DEIB's Young Researchers* *December 2025*
Paper: "Factored-Reward Bandits with Intermediate Observations: Regret Minimization and Best Arm Identification"
Link: `https://www.deib.polimi.it/eng/news/details/1469`

## Personal Research Awards and Recognitions

**Winner of the "Dimitris N. Chorafas Award"** **Chorafas Foundation**
*Two 10,000$ prize awarded to the best two Ph.D. thesis (out of ≈ 100 candidates)* *September 2025*
The Dimitris N. Chorafas Foundation awards scientific prizes for outstanding work in selected fields in the engineering sciences, medicine and the natural sciences. It rewards research with high potential for practical application and special significance attached to its aftermath. The Foundation has 21 partner universities in Europe, North America and Asia (among which MIT, ETHZ, UCLA, Purdue, TUM, EPFL). It currently reserves two annual awards to the Ph.D. in Information Technology programme of Politecnico di Milano.
Link: `https://www.deib.polimi.it/eng/news/details/1463`

**Recipient of a "Springer Award"** **Politecnico di Milano**
*From the Information Technology Ph.D. Board of Professors* *June 2025*
Publication in a PoliMi SpringerBriefs volume, for the best results from the IT Ph.D. program doctors.
Link: `https://link.springer.com/book/10.1007/978-3-032-12359-6`

## Student Awards

**Scholarship Recipient** **Politecnico di Milano**
*For all eligible years of the M.Sc. in Computer Science and Engineering* *2017 – 2019*
Tuition waiver for high academic performance.

# Ph.D. Students Supervision

## Advised Ph.D. Students

[1] Cristiano Migali – Ph.D. Student in Information Technology (XLI cycle) at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. Co-supervisor: Alberto Maria Metelli.

## Co-supervised Ph.D. Students

[1] Alessandro Montenegro – Ph.D. Student in Information Technology (XXXIX cycle) at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. Advisor: Prof. Alberto Maria Metelli. Co-supervisors: Marco Mussi, Matteo Papini and Marcello Restelli. Informal co-supervision.

[2] Gianvito Losapio – Ph.D. Student in Information Technology (XXXIX cycle) at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. Advisor: Prof. Marcello Restelli. Co-supervisors: Alberto Maria Metelli and Marco Mussi. Informal co-supervision.

[3] Federico Corso – Ph.D. Student in Information Technology (XL cycle) at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. Advisor: Prof. Alberto Maria Metelli. Co-supervisors: Marco Mussi and Marcello Restelli. Formal co-supervision.

[4] Diego Alovisetti – Ph.D. Student in Information Technology (XLI cycle) at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. Advisor: Prof. Alberto Maria Metelli. Co-supervisor: Marco Mussi. Formal co-supervision.

# M.Sc. Students Supervision

## Co-supervised M.Sc. Theses

[1] Gianmarco Genalti, "A Multi-Armed Bandit Approach to Dynamic Pricing". Supervisor: Prof. Nicola Gatti. Co-supervisors: Marco Mussi and Alessandro Nuara. M.Sc. in Mathematical Engineering, Politecnico di Milano. December 2021.
Link: `https://hdl.handle.net/10589/183733`

[2] Amedeo Cavallo, "A Combinatorial Multi-Armed Bandit Approach to Online Advertising Budget Optimisation". Supervisor: Prof. Marcello Restelli. Co-supervisors: Marco Mussi and Alessandro Nuara. M.Sc. in Computer Science and Engineering, Politecnico di Milano. December 2021.
Link: Not available due to NDA

[3] Oscar Francesco Pindaro, "Controlling Lithium-Ion Batteries Through Reinforcement Learning". Supervisor: Prof. Marcello Restelli. Co-supervisors: Marco Mussi and Francesco Trovò. M.Sc. in Computer Science and Engineering, Politecnico di Milano. April 2022.
Link: `https://hdl.handle.net/10589/186742`

[4] Davide Lombarda, "Towards Automated Reinforcement Learning". Supervisor: Prof. Marcello Restelli. Co-supervisors: Alberto Maria Metelli, Marco Mussi and Francesco Trovò. M.Sc. in Mathematical Engineering, Politecnico di Milano. April 2022.
Link: `https://hdl.handle.net/10589/187829`

[5] Thomas Petrone, "Hidden Markov Model for Single User Response Prediction in Digital Advertising Campaigns". Supervisor: Prof. Marcello Restelli. Co-supervisor: Marco Mussi. M.Sc. in Mathematical Engineering, Politecnico di Milano. July 2022.
Link: `https://hdl.handle.net/10589/189641`

[6] Alessandro Montenegro, "Best Model Selection via Stochastic Rising Bandits". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Marco Mussi, Marcello Restelli and Francesco Trovò. M.Sc. in Computer Science and Engineering, Politecnico di Milano. May 2023.
Link: `https://hdl.handle.net/10589/210714`

[7] Andrea d'Silva, "Integrating Behavioral Cloning into a Reinforcement Learning pipeline". Supervisor: Prof. Francesco Trovò. Co-supervisors: Alberto Maria Metelli, Marco Mussi and Marcello Restelli. M.Sc. in Computer Science and Engineering, Politecnico di Milano. May 2023.
Link: `https://hdl.handle.net/10589/208354`

[8] Francesco Fulco Gonzales, "Stochastic Linear Bandit with Global-Local Structure". Supervisor: Prof. Francesco Trovò. Co-supervisors: Gianmarco Genalti, Marco Mussi and Marcello Restelli. M.Sc. in

Computer Science and Engineering, Politecnico di Milano. May 2023.
Link: https://hdl.handle.net/10589/210267

[9] Vittorio Arianna, "Multi-Armed Bandits for Joint Pricing and Advertising". Supervisor: Prof. Nicola Gatti. Co-supervisors: Gianmarco Genalti and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. October 2023.
Link: https://hdl.handle.net/10589/211018

[10] Marco Bonalumi, "An Online Learning Algorithm for Real-time Bidding". Supervisor: Prof. Marcello Restelli. Co-supervisors: Gianmarco Genalti and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. December 2023.
Link: https://hdl.handle.net/10589/214882

[11] Alessandro Contù, "Budget Optimization in Marketing Mix Models". Supervisor: Prof. Francesco Trovò. Co-supervisors: Marco Mussi and Marcello Restelli. M.Sc. in Computer Science and Engineering, Politecnico di Milano. December 2023.
Link: https://hdl.handle.net/10589/214238

[12] Andrea Cerasani, "An Online Dynamic Pricing Algorithm for Complementary Products". Supervisor: Prof. Marcello Restelli. Co-supervisors: Alessandro Lavelli and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. December 2023.
Link: https://hdl.handle.net/10589/214318

[13] Federico Corso, "Smoothed OMD: an Algorithm for No-regret Learning in Adversarial MDPs with Revealed Transitions". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Marco Mussi and Riccardo Zamboni. M.Sc. in Automation and Control Engineering, Politecnico di Milano. July 2024.
Link: https://hdl.handle.net/10589/223226

[14] Davide Beretta, "Distributed Reinforcement Learning for Power Grid Operations". Supervisor: Prof. Marcello Restelli. Co-supervisors: Gianvito Losapio, Alberto Maria Metelli and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. October 2024.
Link: https://hdl.handle.net/10589/226817

[15] Valentina Abbattista, "Online Learning for PID Controller Tuning". Supervisor: Prof. Alberto Maria Metelli. Co-supervisor: Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. October 2024.
Link: https://hdl.handle.net/10589/226952

[16] Giacomo Cartechini, "Distributed Reinforcement Learning for Large-Scale Networks". Supervisor: Prof. Marcello Restelli. Co-supervisors: Gianvito Losapio, Marco Mussi and Alberto Maria Metelli. M.Sc. in Computer Science and Engineering, Politecnico di Milano. December 2024.
Link: https://hdl.handle.net/10589/230913

[17] Fabio Patella, "Reinforcement Learning for Digital Advertising Cross-Channel Budget Optimization". Supervisor: Prof. Marcello Restelli. Co-supervisors: Alberto Maria Metelli, Marco Mussi and Alessandro Nuara. M.Sc. in Computer Science and Engineering, Politecnico di Milano. April 2025.
Link: https://hdl.handle.net/10589/234894

[18] Leonardo Cesani, "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, Politecnico di Milano. April 2025.
Link: https://hdl.handle.net/10589/234298

[19] Federico Mansutti, "Trajectory Reuse in Policy Gradients". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Alessandro Montenegro, Matteo Papini, Marco Mussi and Brian Ziebart. M.Sc. in Computer Science and Engineering, Politecnico di Milano. July 2025.
Link: https://hdl.handle.net/10589/239917

[20] Cristiano Migali, "Towards Closing the Gap in Restless Rising Bandits". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Marco Mussi and Gianmarco Genalti. M.Sc. in Computer Science and Engineering, Politecnico di Milano. July 2025.
Link: https://hdl.handle.net/10589/240041

[21] Carlo Fabrizio, "Graph-Based Multi-Agent Reinforcement Learning for Power Grid Control". Supervisor: Prof. Marcello Restelli. Co-supervisors: Marco Mussi, Gianvito Losapio and Alberto Maria Metelli. M.Sc.

in Computer Science and Engineering, Politecnico di Milano. July 2025.
Link: `https://hdl.handle.net/10589/240177`

[22] Leonardo Bianconi, "Multi-Dimensional Reward Learning from Preference Feedback". Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Simone Drago and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. December 2025.
Link: `https://hdl.handle.net/10589/246463`

[23] Andrea Fondacaro, "Dynamic State and Action Factorization for Distributed Reinforcement Learning". Supervisor: Prof. Marcello Restelli. Co-supervisors: Marco Mussi and Gianvito Losapio. M.Sc. in Computer Science and Engineering, Politecnico di Milano. March 2026.
Link: To Appear.

[24] Silvia Firenze. Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Gianmarco Genalti and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. In Progress.

[25] Andrea Zhang. Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Gianmarco Genalti and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. In Progress.

### Co-supervised M.Sc. Projects

[1] Andrea Mastroberti. Multi-disciplinary project on the development of a physical demo of a Reinforcement Learning agent. Project Supervisors: Marcello Restelli and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. 2025.

## Poster Presentations and Oral Talks

- Poster Presentation of "Trading-off Reward Maximization and Stability in Sequential Decision Making" at the 18th European Workshop on Reinforcement Learning (EWRL 2025 – Tübingen, Germany – September 2025)
- Poster Presentation of "A Theoretical Perspective on Sequential Decision Making with Preference Feedback" at the 18th European Workshop on Reinforcement Learning (EWRL 2025 – Tübingen, Germany – September 2025)
- Poster Presentation of "A Novel Self-Normalized Bernstein-Like Dimension-Free Inequality and Regret Bounds for Generalized Kernelized Bandits" at the 18th European Workshop on Reinforcement Learning (EWRL 2025 – Tübingen, Germany – September 2025)
- Poster Presentation of "Gym4ReaL: A Benchmark Suite for Evaluating Reinforcement Learning in Realistic Domains" at the 18th European Workshop on Reinforcement Learning (EWRL 2025 – Tübingen, Germany – September 2025)
- Poster Presentation of "Sleeping Reinforcement Learning" at the 42nd International Conference on Machine Learning (ICML 2025 – Vancouver, B.C., Canada – July 2025)
- Poster Presentation of "Towards Theoretical Understanding of Sequential Decision Making with Preference Feedback" at the 42nd International Conference on Machine Learning (ICML 2025 – Vancouver, B.C., Canada – July 2025)
- Poster Presentation of "Convergence Analysis of Policy Gradient Methods with Dynamic Stochasticity" at the 42nd International Conference on Machine Learning (ICML 2025 – Vancouver, B.C., Canada – July 2025)
- Poster Presentation of "Position: Constants are Critical in Regret Bounds for Reinforcement Learning" at the 42nd International Conference on Machine Learning (ICML 2025 – Vancouver, B.C., Canada – July 2025)
- Poster Presentation of "Factored-Reward Bandits with Intermediate Observations" at the 41st International Conference on Machine Learning (ICML 2024 – Vienna, Austria – July 2024)
- Poster Presentation of "Best Arm Identification for Stochastic Rising Bandits" at the 41st International Conference on Machine Learning (ICML 2024 – Vienna, Austria – July 2024)

- Poster Presentation of "Learning Optimal Deterministic Policies with Stochastic Policy Gradients" at the 41st International Conference on Machine Learning (ICML 2024 – Vienna, Austria – July 2024)
- Poster Presentation of "Graph-Triggered Rising Bandits" at the 41st International Conference on Machine Learning (ICML 2024 – Vienna, Austria – July 2024)
- Poster Presentation of "Open Problem: Tight Bounds for Bernoulli Rewards in Kernelized Multi-Armed Bandits" at the Workshop on Aligning Reinforcement Learning Experimentalists and Theorists, part of the 41st International Conference on Machine Learning (ICML 2024 – Vienna, Austria – July 2024)
- Poster Presentation of "Autoregressive Bandits" at the 27th International Conference on Artificial Intelligence and Statistics (AISTATS 2024 – Valencia, Spain – May 2024)
- Poster Presentation of "Online Learning in Autoregressive Dynamics" at the 16th European Workshop on Reinforcement Learning (EWRL 2023 – Brussels, Belgium – September 2023)
- Poster Presentation of "Stochastic Rising Bandits: A Best Arm Identification Approach" at the 16th European Workshop on Reinforcement Learning (EWRL 2023 – Brussels, Belgium – September 2023)
- Poster Presentation of "Dynamical Linear Bandits" at the 40th International Conference on Machine Learning (ICML 2023 – Honolulu, Hawaii, USA – July 2023)
- Poster Presentation of "A Best Arm Identification Approach for Stochastic Rising Bandits" at the Workshop on New Frontiers in Learning, Control, and Dynamical Systems, part of the 40th International Conference on Machine Learning (ICML 2023 – Honolulu, Hawaii, USA – July 2023)
- Oral Presentation of "Dynamic Pricing with Volume Discounts in Online Settings" at the 35th Conference on Innovative Applications of Artificial Intelligence, part of the 37th AAAI Conference on Artificial Intelligence (IAAI/AAAI 2023 – Virtual – February 2023)
- Poster Presentation of "Dynamic Pricing with Online Data Aggregation and Learning" at the 15th European Workshop on Reinforcement Learning (EWRL 2022 – Milan, Italy – September 2022)
- Poster Presentation of "Pricing the Long Tail by Explainable Product Aggregation and Monotonic Bandits" at the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2022 – Washington D.C., USA – August 2022)
- Poster Presentation of "Dynamical Linear Bandits for Long-Lasting Vanishing Rewards" at the Complex Feedback in Online Learning Workshop, part of the 39th International Conference on Machine Learning (ICML 2022 – Baltimore, Maryland, USA – July 2022)

## Other Academic Activities

### Participation to International Conferences, Workshops, and Summer Schools

- 18th European Workshop on Reinforcement Learning – EWRL 2025
  Tübingen, Germany. September 2025.
- 42nd International Conference on Machine Learning – ICML 2025
  Vancouver, B.C., Canada. July 2025.
- 41st International Conference on Machine Learning – ICML 2024
  Vienna, Austria. July 2024.
- 27th International Conference on Artificial Intelligence and Statistics – AISTATS 2024
  Valencia, Spain. May 2024.
- 16th European Workshop on Reinforcement Learning – EWRL 2023
  Brussels, Belgium. September 2023.
- 40th International Conference on Machine Learning – ICML 2023
  Honolulu, Hawaii, USA. July 2023.
- Reinforcement Learning Summer School – RLSS 2023
  Barcelona, Spain. June 2023.

- 15th European Workshop on Reinforcement Learning – EWRL 2022
  Milan, Italy. September 2022.
- 28th ACM International Conference on Knowledge Discovery and Data Mining – KDD 2022
  Washington D.C., USA. August 2022.
- 39th International Conference on Machine Learning – ICML 2022
  Baltimore, Maryland, USA. July 2022.
- DeepLearn Summer School – DeepLearn 2021
  Virtual. July 2021.

## Area Chair Activities

**International Conferences:**

- Neural Information Processing Systems (NeurIPS) — 2026
- ACM SIGKDD Conference on Knowledge Discovery and Data Mining — 2026

## Reviewer Activities

**International Conferences:**

- Neural Information Processing Systems (NeurIPS) — 2023, 2024, 2025
- International Conference on Machine Learning (ICML) — 2023, 2024, 2025, 2026
- International Conference on Learning Representations (ICLR) — 2024
- International Conference on Artificial Intelligence and Statistics (AISTATS) — 2023, 2026
- AAAI Conference on Artificial Intelligence (AAAI) — 2024, 2025, 2026
- International Conference on Automated Machine Learning (AutoML) — 2022

**International Journals:**

- Springer – Machine Learning (Q1)
- IEEE – Transactions on Neural Networks and Learning Systems (Q1)
- IEEE – Robotics and Automation Letters (Q1)
- Elsevier – Engineering Applications of Artificial Intelligence (Q1)

**International Workshops:**

- European Workshop on Reinforcement Learning (EWRL) — 2023, 2024, 2025, 2026
- AutoRL @ ICML 2024
- ARLET @ ICML 2024

Last update: May 8$^{th}$, 2026