Marco Mussi | Postdoctoral Researcher

marco.mussi@polimi.it • marcomussi.github.io • marcomussi95 • marcomussi
0000-0001-8356-6744

Summary

I am a Postdoctoral Researcher with 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.

Highlights

- Currently **Postdoctoral Researcher** 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 **5 publications in peer-reviewed international journals** (all Q1 and all as main contributor). Author of **9 publications in peer-reviewed international conferences** (6 as main contributor), including ICML (5), NeurIPS (1), KDD (1), and AISTATS (1). According to Core Conference Rating: 7 publications in A* venues and 1 publication in A venues. Author of 8 publications in peer-reviewed international workshops.
- **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%).
- Co-supervisor of **3 Ph.D. students** (all ongoing) and **23 M.Sc. students theses** (5 ongoing).
- **Teaching Assistant** of the B.Sc. course of Computer Science (Civil Engineering, delivered in english) and Laboratory Teaching Assistant of the B.Sc. course of Computer Science (Mechanical Engineering) at Politecnico di Milano.
- **Co-organizer** of the 15th **European Workshop on Reinforcement Learning** (EWRL 2022).
- Participation in **5 industrial research** projects funded by private companies and public research institutions.
- Participation in the **Horizon Europe** competitive research project **AI4REALNET**.

Education

Ph.D. in Information Technology (Dottorato di Ricerca) — **Milano**
Politecnico di Milano – Dipartimento di Elettronica, Informazione e Bioingegneria — *Nov 2020 – Jun 2024*
Focus on Reinforcement Learning and Online Learning
Supervisor: Prof. Marcello Restelli
Final Mark: Laude

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
Scholarship: Tuition waiver for high academic performance
Master thesis: Improving Aerodynamic Load Estimation Algorithms for F1 Racing Cars
Supervisor: Prof. Marcello Restelli
Industrial Partner: Scuderia Ferrari F1

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

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

Professional Experience

Research Scientist **Milano**
ML cube *Nov 2020 – Jun 2024*
Goal: develop algorithms for dynamic pricing and advertising optimization

Academic Experience

Postdoctoral Researcher (Assegno di Ricerca) **Milano**
Politecnico di Milano *Jun 2024 – now*
Supervisor: Prof. Marcello Restelli

Research Assistant (Assegno di Ricerca) **Milano**
Politecnico di Milano *Jan 2020 – Oct 2020*
Supervisor: Prof. Marcello Restelli

Competitive Research Projects

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
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.

Industrial Projects

MLcube (Politecnico di Milano 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 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 a Politecnico di Milano's spin-off MLcube *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
Funding: Winner of the ELISE's 2nd Open Call Grants

Dynamic Pricing for E-commerce **Milano**
Collaboration with an e-commerce website *2021 – 2022*
Focus: Development and release of a dynamic pricing solution for an e-commerce website managing over 20000 products

Open Source Software Release

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

Teaching

Teaching Assistant **24 Hours**
Course of Foundations of 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. ∼150 students. 6 CFU.
Lecturer: Prof. Alberto Maria Metelli.
Course delivered in English.
Students' Evaluation: Pending.

Laboratory Teaching Assistant **15 Hours**
Course of 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: Pending.

Teaching Assistant **24 Hours**
Course of Foundations of 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. ∼100 students. 6 CFU.
Lecturer: Prof. Alberto Maria Metelli.
Course delivered in English.
Students' Evaluation: 3.4/4 (University Average: 3.1/4).

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

Organization of International Events

AI for Safety-Critical Infrastructures AI-SCI Workshop at ECML-PKDD 2025 **Porto**
Organizer and Program Chair
To happen. Will be held within the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD).

15th European Workshop on Reinforcement Learning (EWRL 2022) **Milano**
Organizer and Communication Chair

Held in the Bovisa Campus of Politecnico di Milano from 19 to 21 September 2022.
Attendance: ∼150. Number of submitted papers: ∼100.

Seminars

AI4REALNET Dissemination Webinar — **AI4REALNET Consortium**
Distributed and Hierarchical Reinforcement Learning (Invited) — *24 April 2024*

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

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

Theses

M.Sc. Thesis

Title: *Improving Aerodynamic Load Estimation Algorithms for F1 Racing Cars*

Supervisor: Prof. Marcello Restelli

Industrial Partner: Scuderia Ferrari F1

Short Abstract: The thesis studies the aerodynamic behavior of Formula One cars, aiming to develop effective methodologies for the estimation of aerodynamic forces on the vehicle. Using data-driven techniques coming from the Machine Learning field that exploits the data gathered during the wind tunnel tests, and the measurements from a small set of pressure sensors, this work allows to reconstruct the pressure profile of the entire car. Moreover, this work proposes a method to aggregate the data coming from a set of wind tunnel test, to better generalize the aerodynamic load estimation over newly seen aerodynamic configurations.

Ph.D. Thesis

Title: *Online Learning Methods for Pricing and Advertising*

Supervisor: Prof. Marcello Restelli

Short Abstract: Nowadays, when it comes to selling a product online, two of the most significant factors are the pricing strategy and the investments in advertising. When determining the price of a product, it is essential to strike a balance. The price should neither be set too low, as this would result in a reduced revenue from the single sale, nor too high, as it may deter potential buyers. The amount of money we invest in advertising should be balanced to let people know our offer without overspending or reaching people who are not interested. These two aspects are usually handled disjointedly by humans, but this, even if we proceed to optimize for the two components individually, may lead to a suboptimal solution. This work focuses on the adoption of online learning to solve the task of finding the optimal price for a product and how to advertise it properly, offering both theoretical frameworks and practical solutions for addressing the associated challenges.

Conference Awards and Recognitions

[A1] Spotlight Paper at ICML 2024 for "Learning Optimal Deterministic Policies with Stochastic Policy Gradients" – top 3.5% over 9473 submissions (Core 2023: A* – GGS 2021: A++)

[A2] Spotlight Paper at ICML 2024 for "Best Arm Identification for Stochastic Rising Bandits" – top 3.5% over 9473 submissions (Core 2023: A* – GGS 2021: A++)

[A3] Oral Presentation at KDD 2022 for "Pricing the Long Tail by Explainable Product Aggregation and Monotonic Bandits" – top 7% over 753 submissions (Core 2021: A* – GGS 2021: A++)

List of Publications

Refereed International Conferences Papers

[C1] 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`

[C2] 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

[C3] 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

[C4] 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

[C5] 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

[C6] 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

[C7] 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

[C8] Marco Mussi*, Gianmarco Genalti*, Alessandro Nuara, Francesco Trovó, Nicola Gatti and Marcello Restelli. Dynamic Pricing with Volume Discounts in Online Settings. Proceedings of the Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence (IAAI), volume 37, pages 15560–15568, AAAI, 2023. Innovative Application of AI Award.
Link: https://doi.org/10.1609/aaai.v37i13.26845

[C9] 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 Journal Articles

[J1] 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
Link: https://doi.org/10.1613/jair.1.17614

[J2] 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
Link: https://doi.org/10.1016/j.est.2024.110572

[J3] 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
Link: https://doi.org/10.1016/j.eswa.2023.119883

[J4] 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
Link: https://doi.org/10.1016/j.est.2022.105467

[J5] 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
Link: https://doi.org/10.1016/j.est.2021.103309

Refereed International Workshop Papers

[W1] 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.
Link: https://marcomussi.github.io/papers/ecmlw2024/paper.pdf

[W2] 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

[W3] 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

[W4] 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

[W5] 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

[W6] 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

[W7] 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 – 10/96)
Link: https://marcomussi.github.io/papers/ewrl2022/paper.pdf

[W8] 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

[T1] 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

Preprints

[P1] 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

[P2] 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

[P3] 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

[R1] 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.

[R2] Marco Mussi*, Simone Drago*, Marcello Restelli and Alberto Maria Metelli. Factored-Reward Bandits with Intermediate Observations: Regret Minimization and Best Arm Identification. 2024. Under Review at the Artificial Intelligence Journal.

[R3] 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.

[R4] Simone Drago, Marco Mussi and Alberto Maria Metelli. Position: Constants are Critical in Regret Bounds for Reinforcement Learning. 2025. Under Review at the International Conference on Machine Learning — Position Paper Track (ICML 2025).

[R5] Simone Drago*, Marco Mussi* and Alberto Maria Metelli. Sleeping Reinforcement Learning. 2025. Under Review at the International Conference on Machine Learning (ICML 2025).

[R6] Simone Drago, Marco Mussi and Alberto Maria Metelli. Towards Theoretical Understanding of Sequential Decision Making with Preference Feedback. 2025. Under Review at the International Conference on Machine Learning (ICML 2025).

[R7] Gianmarco Tedeschi, Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Trading-off Statistical and Computational Efficiency via W-step Markov Decision Process: A Policy Gradient Approach. 2025. Under Review at the International Conference on Machine Learning (ICML 2025).

[R8] Alessandro Montenegro, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Convergence Analysis of Policy Gradient Methods with Dynamic Stochasticity. 2025. Under Review at the International Conference on Machine Learning (ICML 2025).

[R9] Alberto Maria Metelli, Simone Drago and Marco Mussi. Generalized Kernelized Bandits: Novel Self-normalized Bernstein-Like Dimension-Free Inequality and Regret Bounds. 2025. Under Review at the Annual Conference on Learning Theory (COLT 2025).

[R10] Marco Mussi, Alberto Maria Metelli, Marcello Restelli, Gianvito Losapio, Ricardo Jorge Bessa, Daniel Boos, Clark Borst, 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, Giacomo Zanotti. Human-AI Interaction in Safety-Critical Network Infrastructures. 2025. Under Review at iScience (Cell Press).

In Preparation

[P1] Marco Mussi, Andrea Cerasani, Alessandro Lavelli and Marcello Restelli. Online Dynamic Pricing of Complementary Goods.

[P2] Alessandro Montenegro, Leonardo Cesani, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Deterministic Policy Gradients with Constraints.

[P3] Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Gaussian Processes for Demand Learning in Pricing.

[P4] Marco Mussi, Marcello Restelli and Alberto Maria Metelli. Behavioral Cloning from Human Feedback.

[P5] Simone Drago, Marco Mussi and Alberto Maria Metelli. Learning in Preference-based Markov Decision Processes.

[P6] Federico Corso, Marco Mussi and Alberto Maria Metelli. Learning in Markov Decision Processes with Stability Constraints.

[P7] Cristiano Migali, Marco Mussi, Gianmarco Genalti and Alberto Maria Metelli. Tight Bounds for Restless Rising Bandits.

[P8] Alessandro Montenegro, Marco Mussi, Federico Mansutti, Matteo Papini and Alberto Maria Metelli. On Sample Reuse for Policy Gradients.

Ph.D. Students Co-supervision

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

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

M.Sc. Students Co-supervision

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-supervisor: Alberto Maria Metelli, Marco Mussi and Alessandro Nuara. M.Sc. in Computer Science and Engineering, Politecnico di Milano. April 2025.
Link: To Appear

[18] Leonardo Cesani, "Learning Deterministic Policies in Constrained Markov Decision Processes with Policy Gradients". Supervisor: Prof. Matteo Papini. Co-supervisor: Alberto Maria Metelli, Alessandro Montenegro and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. April 2025.
Link: To Appear

[19] Federico Mansutti. Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Alessandro Montenegro, Marco Mussi and Matteo Papini. M.Sc. in Computer Science and Engineering, Politecnico di Milano. In Progress.

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

[21] Carlo Fabrizio. Supervisor: Prof. Marcello Restelli. Co-supervisors: Gianvito Losapio, Marco Mussi and Alberto Maria Metelli. M.Sc. in Computer Science and Engineering, Politecnico di Milano. In Progress.

[22] Leonardo Bianconi. Supervisor: Prof. Alberto Maria Metelli. Co-supervisors: Simone Drago and Marco Mussi. M.Sc. in Computer Science and Engineering, Politecnico di Milano. In Progress.

[23] Andrea Fondacaro. Supervisor: Prof. Marcello Restelli. Co-supervisors: Gianvito Losapio, Marco Mussi and Alberto Maria Metelli. M.Sc. in Computer Science and Engineering, Politecnico di Milano. In Progress.

Poster and Oral Talks

- 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 (ARLET @ 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 "Open Problem: Tight Bounds for Bernoulli Rewards in Kernelized Multi-Armed Bandits" at the Complex Feedback in Online Learning Workshop, part of the 39th International Conference on Machine Learning (CFOL @ ICML 2022 – Baltimore, Maryland, USA – July 2022)

Other Academic Activities

Participation to International Conferences, Workshops, and Summer Schools

- International Conference on Machine Learning – ICML 2024
  Vienna, Austria. July 2024.
- International Conference on Artificial Intelligence and Statistics – AISTATS 2024
  Valencia, Spain. May 2024.
- European Workshop on Reinforcement Learning – EWRL 2023
  Brussels, Belgium. September 2023.
- International Conference on Machine Learning – ICML 2023
  Honolulu, Hawaii, USA. July 2023.
- Reinforcement Learning Summer School – RLSS 2023
  Barcelona, Spain. June 2023.
- European Workshop on Reinforcement Learning – EWRL 2022
  Milan, Italy. September 2022.
- ACM International Conference on Knowledge Discovery and Data Mining – KDD 2022
  Washington D.C., USA. August 2022.
- International Conference on Machine Learning – ICML 2022
  Baltimore, Maryland, USA. July 2022.
- DeepLearn Summer School – DeepLearn 2021
  Virtual. July 2021.

Reviewer Activities

Reviewer for International Conferences:

- Neural Information Processing Systems (NeurIPS)
- International Conference on Machine Learning (ICML)
- International Conference on Learning Representations (ICLR)
- International Conference on Artificial Intelligence and Statistics (AISTATS)
- AAAI Conference on Artificial Intelligence (AAAI)
- International Conference on Automated Machine Learning (AutoML)

Reviewer for 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)

Reviewer for International Workshops:

- European Workshop on Reinforcement Learning (EWRL)
- AutoRL @ ICML 2024
- ARLET @ ICML 2024

Last update: April 4th, 2025