

Marco Mussi | Postdoctoral Researcher

✉ marco.mussi@polimi.it • 🌐 marcomussi.github.io • in marcomussi95
📍 marcomussi • 📞 0000-0001-8356-6744

Short Bio

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 and public Italian companies.

Education

Ph.D. in Information Technology <i>Politecnico di Milano</i> Focus on Reinforcement Learning and Online Learning. <i>Supervisor:</i> Prof. Marcello Restelli <i>Co-supervisor:</i> Prof. Alberto Maria Metelli <i>Tutor:</i> Prof. Nicola Gatti <i>Industrial Partner:</i> ML cube	Milano Nov 2020 – Jun 2024
M.Sc. in Computer Science and Engineering <i>Politecnico di Milano</i> <i>Main focus:</i> Artificial Intelligence and Machine Learning <i>Scholarship:</i> Tuition waiver for high academic performance <i>Master thesis:</i> <i>Improving Aerodynamic Load Estimation Algorithms for F1 Racing Cars</i> <i>Supervisor:</i> Prof. Marcello Restelli <i>Industrial Partner:</i> Scuderia Ferrari F1	Milano Sep 2017 – Dec 2019
B.Sc. in Computer Science and Engineering <i>Politecnico di Milano</i>	Milano Sep 2014 – Jul 2017
High School Diploma in Computer Science <i>IIS Galileo Galilei</i>	Crema Sep 2008 – Jul 2014

Experience

Academic

Politecnico di Milano <i>Postdoctoral Researcher</i> <i>Supervisor:</i> Prof. Marcello Restelli	Milano Jun 2024 – now
Politecnico di Milano <i>Research Assistant</i> <i>Supervisor:</i> Prof. Marcello Restelli	Milano Jan 2020 – Oct 2020

Professional

ML cube <i>Research Scientist</i> <i>Goal:</i> develop algorithms for dynamic pricing and advertising optimization	Milano Nov 2020 – now
---------------------------------------------------------------------------------------------------------------------------------	---------------------------------

Languages

Italian: Mother Tongue

English: Excellent

Industrial Projects

AD cube Marketing Mix Model

Collaboration with a Politecnico di Milano's spin-off

Focus: Budget optimization in advertising, considering advertising campaigns interactions

Milano

Nov 2022 – Jun 2024

Data-driven Optimization Marketing Mix Models for Advertising

Collaboration with an Italian marketing agency

Focus: Implementation of a MMM to solve the attribution problem in digital advertising in contexts with scarce and noisy data

Milano

Feb 2022 – Aug 2022

Dynamic Pricing for E-commerce

Collaboration with an Italian e-commerce

Focus: Implementation of a dynamic pricing model for an e-commerce website managing over 20000 products

Milano

Feb 2021 – May 2022

AD cube Product Release

Collaboration with a Politecnico di Milano's spin-off

Focus: Release of AD cube, a product for advertising optimization in online campaigns

Milano

Nov 2020 – Feb 2022

Last-mile Delivery Optimization

Collaboration with a last-mile delivery company

Focus: Delivery allocation using Reinforcement Learning and bikers load estimation using Supervised Learning techniques

Milano

May 2020 – Oct 2020

Reinforcement Learning in Smart-grids

Collaboration with an Energy Research Center

Focus: Exploit Reinforcement Learning solutions to preserve the battery State of Health in smart-grids

Milano

Feb 2020 – Oct 2022

European Projects

AI4REALNET

Part of the team for the development of foundational algorithms

Focus: The scope of AI4REALNET covers the perspective of AI-based solutions addressing critical systems (electricity, railway, and air traffic management) modeled by networks that can be simulated, and are traditionally operated by humans, and where AI systems complement and augment human abilities. It has two main strategic goals: (i) to develop the next generation of decision-making methods powered by supervised and reinforcement learning, which aim at trustworthiness in AI-assisted human control with augmented cognition, hybrid human-AI co-learning and autonomous AI, with the resilience, safety, and security of critical infrastructures as core requirements, and (ii) to boost the development and validation of novel AI algorithms, by the consortium and AI community, through existing open-source digital environments capable of emulating realistic scenarios of physical systems operation and human decision-making.

Politecnico di Milano

Oct 2023 – now

Master 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.

Teaching

- Exercise Sessions Lecturer - Politecnico di Milano** **24 Hours**
Lecturer for the Exercise Sessions of the Course of Foundations of Computer Science Feb 2024 – Jun 2024
Exercise sessions mainly on C and Fortran programming languages.
Course delivered in English.
- Tutor - Politecnico di Milano and Cefriel** **30 Hours**
Tutor for a Master in AI/ML Sep 2022 – Jul 2023
Supervision of a team in the application of Reinforcement Learning algorithms to real-world control problems.

Organization of International Events

- European Workshop on Reinforcement Learning** **Milano**
Communication Chair
Communication Chair for the 15th edition of the European Workshop on Reinforcement Learning (EWRL) 2022

Seminars

- AI4REALNET Dissemination Webinar** **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** **Ricerca Sistema Energetico**
A Data-Driven Method for State of Charge Estimation 23 October 2020

Publications

International Conferences

- [C1] [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), 2024. (A* Core Ranking - 2609/9473, Acceptance rate 27.5%)
- [C2] [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), 2024. (A* Core Ranking - Spotlight - 335/9473, top 3.5%)
- [C3] 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), 2024. (A* Core Ranking - Spotlight - 335/9473, top 3.5%)
- [C4] 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), 2024. (A* Core Ranking - 2609/9473, Acceptance rate 27.5%)
- [C5] 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), 2024. (A Core Ranking - 546/1980, Acceptance rate 27.6%)
- [C6] [Marco Mussi](#), Alberto Maria Metelli and Marcello Restelli. Dynamical Linear Bandits. Proceedings of the 40th International Conference on Machine Learning (ICML), 2023. (A* Core Ranking - 1827/6538, Acceptance rate 27.9%)
- [C7] [Marco Mussi](#), Gianmarco Genalti, Alessandro Nuara, Francesco Trovò, 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), 2023. AAAI. Innovative Application of AI Award.
- [C8] [Marco Mussi](#)^{*}, Gianmarco Genalti^{*}, Francesco Trovò, 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, 2022. (A* Core Ranking - Oral Presentation - 54/753, top 7%)

Journals

[J1] 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*, 82, 2024. (Q1 Scimago)

[J2] Marco Mussi, Davide Lombarda, Alberto Maria Metelli, Francesco Trovó and Marcello Restelli. ARLO: A Framework for Automated Reinforcement Learning. *Expert Systems with Applications*, 224, 2023. (Q1 Scimago)

[J3] 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*, 55, 2022. (Q1 Scimago)

[J4] 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*, 44, 2021. (Q1 Scimago)

Workshops

[W1] 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.

[W2] 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.

[W3] 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.

[W4] Alessandro Montenegro, Marco Mussi, Francesco Trovó, Marcello Restelli and Alberto Maria Metelli. Stochastic Rising Bandits: A Best Arm Identification Approach. *European Workshop on Reinforcement Learning (EWRL)*. 2023.

[W5] Alessandro Montenegro, Marco Mussi, Francesco Trovó, Marcello Restelli and Alberto Maria Metelli. A Best Arm Identification Approach for Stochastic Rising Bandits. *Workshop on New Frontiers in Learning, Control, and Dynamical Systems at International Conference of Machine Learning (ICML)*. 2023.

[W6] 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)

[W7] Marco Mussi, Alberto Maria Metelli and Marcello Restelli. Dynamical Linear Bandits for Long-Lasting Vanishing Rewards. *Complex Feedback in Online Learning Workshop at International Conference of Machine Learning (ICML)*. 2022.

Under Review

[R1] Marco Mussi and Alberto Maria Metelli. Generalizing the Regret: an Analysis of Lower and Upper Bounds. 2024. Submitted to the *Journal of Artificial Intelligence Research (JAIR)*.

[R2] Simone Drago*, Marco Mussi* and Alberto Maria Metelli. Sleeping Reinforcement Learning. 2024. Submitted to *Neural Information Processing Systems (NeurIPS)*.

[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. Submitted to *Neural Information Processing Systems (NeurIPS)*.

[R4] Alessandro Montenegro, Marco Mussi, Matteo Papini and Alberto Maria Metelli. Last-Iterate Global Convergence of Policy Gradients for Constrained Reinforcement Learning. 2024. Submitted to *Neural Information Processing Systems (NeurIPS)*.

[R5] Gianvito Losapio, Davide Beretta, Marco Mussi, Alberto Maria Metelli and Marcello Restelli. State and Action Factorization in Power Grids. Submitted to the *Workshop on Machine Learning for Sustainable Power Systems at the European Conference on Machine Learning (ECML)*.

In Preparation.....

- [P1] [Marco Mussi](#)^{*}, [Andrea Cerasani](#)^{*}, [Alessandro Lavelli](#) and [Marcello Restelli](#). Online Dynamic Pricing of Complementary Goods.
- [P2] [Marco Mussi](#)^{*}, [Simone Drago](#)^{*}, [Marcello Restelli](#) and [Alberto Maria Metelli](#). Factored-Reward Bandits with Intermediate Observations: Regret Minimization and Best Arm Identification.
- [P3] [Simone Drago](#), [Marco Mussi](#) and [Alberto Maria Metelli](#). Notes on Minimax Regret Bounds for Reinforcement Learning.
- [P4] [Marco Mussi](#), [Simone Drago](#) and [Alberto Maria Metelli](#). Open Problem: Tight Bounds for Kernelized Multi-Armed Bandits with Bernoulli Rewards.
- [P5] [Gianmarco Genalti](#), [Marco Mussi](#), [Nicola Gatti](#), [Marcello Restelli](#), [Matteo Castiglioni](#) and [Alberto Maria Metelli](#). Graph-Triggered Rising and Rotting Bandits.
- [P6] [Alberto Maria Metelli](#), [Valentina Abbattista](#) and [Marco Mussi](#). Online Learning for PID Controllers.
- [P7] AI4REALNET Consortium. AI for the Operation of Critical Infrastructures. Position/Consortium Paper.
- [P8] [Alessandro Montenegro](#), [Leonardo Cesani](#), [Marco Mussi](#), [Matteo Papini](#) and [Alberto Maria Metelli](#). Deterministic Policy Gradients with Constraints.
- [P9] [Alessandro Montenegro](#), [Marco Mussi](#), [Matteo Papini](#) and [Alberto Maria Metelli](#). Deterministic Policy Gradients with Adaptive Variance.
- [P10] [Alessandro Montenegro](#), [Federico Mansutti](#), [Marco Mussi](#), [Matteo Papini](#) and [Alberto Maria Metelli](#). Sample Reuse for Policy Gradients.

Master's Students Supervision

- [1] [Gianmarco Genalti](#) - "A Multi-Armed Bandit Approach to Dynamic Pricing". Co-supervision. Supervisor: Prof. [Nicola Gatti](#) (M.Sc. in Mathematical Engineering, Dec 2021)
- [2] [Amedeo Cavallo](#) - "A Combinatorial Multi-Armed Bandit Approach to Online Advertising Budget Optimisation". Co-supervision. Supervisor: Prof. [Marcello Restelli](#) (M.Sc. in Computer Science and Engineering, Dec 2021)
- [3] [Oscar Francesco Pindaro](#) - "Controlling Lithium-Ion Batteries Through Reinforcement Learning". Co-supervision. Supervisor: Prof. [Marcello Restelli](#) (M.Sc. in Computer Science and Engineering, Apr 2022)
- [4] [Davide Lombarda](#) - "Towards Automated Reinforcement Learning". Co-supervision. Supervisor: Prof. [Marcello Restelli](#) (M.Sc. in Mathematical Engineering, Apr 2022)
- [5] [Thomas Petrone](#) - "Hidden Markov Model for Single User Response Prediction in Digital Advertising Campaigns". Co-supervision. Supervisor: Prof. [Marcello Restelli](#) (M.Sc. in Mathematical Engineering, Jul 2022)
- [6] [Alessandro Montenegro](#) - "Best Model Selection via Stochastic Rising Bandits". Co-supervision. Supervisor: Prof. [Alberto Maria Metelli](#) (M.Sc. in Computer Science and Engineering, May 2023)
- [7] [Andrea d'Silva](#) - "Integrating Behavioral Cloning into a Reinforcement Learning pipeline". Co-supervision. Supervisor: Prof. [Francesco Trovò](#) (M.Sc. in Computer Science and Engineering, May 2023)
- [8] [Francesco Gonzales](#) - "Stochastic Linear Bandit with Global-Local Structure". Co-supervision. Supervisor: Prof. [Francesco Trovò](#) (M.Sc. in Computer Science and Engineering, May 2023)
- [9] [Vittorio Arianna](#) - "Multi-Armed Bandits for Joint Pricing and Advertising". Co-supervision. Supervisor: Prof. [Nicola Gatti](#) (M.Sc. in Computer Science and Engineering, Oct 2023)
- [10] [Marco Bonalumi](#) - "An Online Learning Algorithm for Real-time Bidding". Co-supervision. Supervisor: Prof. [Marcello Restelli](#) (M.Sc. in Computer Science and Engineering, Dec 2023)
- [11] [Alessandro Contù](#) - "Budget Optimization in Marketing Mix Models". Co-supervision. Supervisor: Prof. [Francesco Trovò](#) (M.Sc. in Computer Science and Engineering, Dec 2023)
- [12] [Andrea Cerasani](#) - "An Online Dynamic Pricing Algorithm for Complementary Products". Co-supervision. Supervisor: Prof. [Marcello Restelli](#) (M.Sc. in Computer Science and Engineering, Dec 2023)
- [13] [Federico Corso](#) - "Smoothed OMD: an Algorithm for No-regret Learning in Adversarial MDPs with Revealed Transitions". Co-supervision. Supervisor: Prof. [Alberto Maria Metelli](#) (M.Sc. in Automation and Control Engineering, Jul 2024)

- [14] Valentina Abbattista. Co-supervision. (M.Sc. in Computer Science and Engineering, in progress)
- [15] Federico Mansutti. Co-supervision. (M.Sc. in Computer Science and Engineering, in progress)
- [16] Davide Beretta. Co-supervision. (M.Sc. in Computer Science and Engineering, in progress)
- [17] Fabio Patella. Co-supervision. (M.Sc. in Computer Science and Engineering, in progress)
- [18] Giacomo Cartechini. Co-supervision. (M.Sc. in Computer Science and Engineering, in progress)
- [19] Leonardo Cesani. Co-supervision. (M.Sc. in Computer Science and Engineering, in progress)

Other Academic Activities

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)
- 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

Participation to International Conferences and Workshops

- 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.
- 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.

Participation to Summer Schools

- Reinforcement Learning Summer School - RLSS 2023
Barcelona, Spain. June 2023.

- DeepLearn Summer School - DeepLearn 2021
Virtual. July 2021.

Last update: June 29th, 2024

Autorizzo al trattamento dati ai sensi del GDPR 2016/679 del 27 aprile 2016 (Regolamento Europeo relativo alla protezione delle persone fisiche per quanto riguarda il trattamento dei dati personali).

Il sottoscritto, Marco Mussi, consapevole delle sanzioni penali previste dall'art. 76 T.U. N. 445 del 28.12.2000 sulle disposizioni legislative e regolamentari in materia di documentazione amministrativa, nel caso di mendaci dichiarazioni, o esibizione di atti falsi o contenenti dati non rispondenti a verità, dichiara sotto la propria responsabilità che quanto sopra riportato risponde a verità.

Autorizzo la pubblicazione del Curriculum Vitae sul sito istituzionale del Politecnico di Milano (sez. Amministrazione Trasparente) in ottemperanza al D. Lgs n. 33 del 14 marzo 2013 (e s.m.i.).