

UNIVERSITÀ DEGLI STUDI DI MILANO**selezione pubblica per n. 2 posti di Ricercatore a tempo determinato in tenure track (RTT)****per il settore concorsuale 01/B1 - Informatica ,
settore scientifico-disciplinare INF/01 - Informatica****presso il Dipartimento di INFORMATICA "GIOVANNI DEGLI ANTONI",****(avviso bando pubblicato sulla G.U. n. 41 del 21/05/2024) Codice concorso 5551**

Matteo Papini

CURRICULUM VITAE

INFORMAZIONI PERSONALI

COGNOME	PAPINI
NOME	MATTEO
DATA DI NASCITA	[REDACTED]
[REDACTED]	[REDACTED]
INDIRIZZO	DIPARTIMENTO DI ELETTRONICA, INFORMAZIONE E BIOINGEGNERIA, POLITECNICO DI MILANO, VIA CAMILLO GOLGI 39 (EDIFICIO 21), 20133 MILANO

HIGHLIGHTS

Matteo Papini è ricercatore (**RTD-A**) al **Politecnico di Milano**, presso il Laboratorio di Intelligenza Artificiale e Robotica del Dipartimento di Ingegneria Industriale e dell'Informazione. La sua ricerca è nel campo dell'**intelligenza artificiale** e del **machine learning**, con particolare attenzione al problema delle decisioni sequenziali in condizioni di incertezza, utilizzando la teoria e soluzioni algoritmiche dell'apprendimento per rinforzo (**reinforcement learning**).

Nel 2021 ha conseguito il dottorato con lode in Information Technology presso il Politecnico di Milano, sotto la supervisione di Marcello Restelli, con una tesi dal titolo *Safe Policy Optimization*. Nel 2020 ha lavorato come research intern presso **Facebook AI Research** (ora Meta). Dal 2021 al 2023 è stato ricercatore post-dottorato nel gruppo di ricerca di Artificial Intelligence and Machine Learning presso **Universitat Pompeu Fabra** (UPF), Barcellona, Spagna, nel gruppo di ricerca di Gergely Neu.

È autore di oltre venti articoli peer-reviewed presentati in conferenze di alto livello nel campo dell'intelligenza artificiale e del machine learning, come **NeurIPS** (5 paper), **ICML** (5), **COLT** (2), **AAAI** (2) e **IJCAI** (2). Uno di questi ha ricevuto una presentazione orale a NeurIPS 2018, concessa solo al top 3% dei paper accettati. Ha inoltre pubblicato articoli su riviste di machine learning peer-reviewed (JMLR e Machine Learning). Alcuni dei suoi co-autori ricorrenti sono Marcello Restelli (Politecnico di Milano), Matteo Pirota, Alessandro Lazaric (Meta) e Gergely Neu (UPF).

È membro di **ELLIS** (European Laboratory for Learning and Intelligent Systems) e ha ricoperto il ruolo di esercitatore per diversi corsi al Politecnico di Milano. Ha partecipato come reviewer nel program committee di varie conferenze, tra cui NeurIPS, ICML e AAAI.

Nel 2023, insieme a Vincent Adam (UPF), ha coordinato l'organizzazione della **Reinforcement Learning Summer School** (RLSS 2023) presso UPF. Sarà local chair di **ALT 2025** (International Conference on Algorithmic Learning Theory) a Milano.

TITOLI

TITOLO DI STUDIO

LAUREA MAGISTRALE IN COMPUTER SCIENCE AND ENGINEERING - INGEGNERIA INFORMATICA, Politecnico di Milano, 27/07/2017, 110/110 con lode

Laurea in Ingegneria Informatica, Politecnico di Milano, 24/07/2015, 110/110 con lode

TITOLO DI DOTTORE DI RICERCA O EQUIVALENTI, OVVERO, PER I SETTORI INTERESSATI, DEL DIPLOMA DI SPECIALIZZAZIONE MEDICA O EQUIVALENTE, CONSEGUITO IN ITALIA O ALL'ESTERO

DOTTORATO IN INGEGNERIA DELL'INFORMAZIONE / INFORMATION TECHNOLOGY, Politecnico di Milano, 11/03/2021, con lode

CONTRATTI DI RICERCA, ASSEGNI DI RICERCA O EQUIVALENTI

- Dal 1 settembre 2023 ad oggi: Ricercatore a tempo determinato di tipo A "Junior" (RTDA) presso Politecnico di Milano, Dipartimento di Elettronica, Informazione e Bioingegneria

- Dall'8 giugno 2021 al 31 agosto 2023, ricercatore post dottorato presso Universitat Pompeu Fabra, Barcellona, Spagna, con i seguenti contratti:

08/06/2021 – 26/07/2021: Investigador de Projectes

27/07/2021 – 21/09/2021: Investigador de Projectes

22/09/2021 – 30/09/2021: Investigador de Projectes

01/10/2021 – 30/11/2022: Investigador de Projectes

01/12/2022 – 31/08/2023: Investigador JCierva- Formació

- Dal 1 aprile 2021 al 31 maggio 2021: attività di supporto alla ricerca presso il Dipartimento di Elettronica, Informazione e Bioingegneria del Politecnico di Milano come incaricato di lavoro autonomo per prestazione occasionale d'opera "Reinforcement Learning Techniques for Developing Artificial Test Drivers on a F1 Simulator"

- Dal 1 settembre 2020 al 18 dicembre 2020: research intern presso Facebook AI Research

- Dal 16 novembre 2017 all'11 marzo 2021: assegno di ricerca presso Dipartimento di Elettronica, Informazione e Bioingegneria del Politecnico di Milano, "Studio e Sviluppo di Tecniche di Apprendimento per Rinforzo in Ambito Industriale"

ATTIVITÀ DIDATTICA A LIVELLO UNIVERSITARIO IN ITALIA O ALL'ESTERO

- Il semestre 2023/2024: Attività di didattica integrativa (esercitatore) per il corso di Informatica, AA 2023/2024, Politecnico di Milano, corso di laurea in Ingegneria Civile, 24 ore

- Il semestre 2023/2024: Attività di didattica integrativa (esercitatore) per il corso di Machine Learning, AA 2023/2024, Politecnico di Milano, corso di laurea magistrale in Computer Science and Engineering, 20 ore

- I semestre 2023/2024: Attività di didattica integrativa (esercitatore) per il corso di Foundations of Artificial Intelligence, AA 2023/2024, Politecnico di Milano, corso di laurea magistrale in Computer Science and Engineering, 16 ore
- II semestre 2021/2022: Attività di didattica integrativa (esercitatore) per il corso di Intelligenza Artificiale, AA 2021/2022, Politecnico di Milano, corso di laurea in Ingegneria Informatica (online), 6 ore
- II semestre 2020/2021: Attività di didattica integrativa (esercitatore) per il corso di Intelligenza Artificiale, AA 2020/2021, Politecnico di Milano, corso di laurea in Ingegneria Informatica (online), 6 ore
- II semestre 2019/2020: Attività di didattica integrativa (esercitatore) per il corso di Intelligenza Artificiale, AA 2019/2020, Politecnico di Milano, corso di laurea in Ingegneria Informatica (online), 6 ore
- I semestre 2019/2020: Attività di didattica integrativa (esercitatore) per il corso di Informatica B, AA 2019/2020, Politecnico di Milano, corso di laurea in Ingegneria Meccanica, 26 ore
- I semestre 2018/2019: Attività di didattica integrativa (esercitatore) per il corso di Informatica B, AA 2018/2019, Politecnico di Milano, corso di laurea in Ingegneria Meccanica, 28 ore
- II semestre 2017/2018: Attività di didattica integrativa (esercitatore) per il corso di Web and Internet Economics, AA 2017/2018, Politecnico di Milano (polo territoriale di Como), corso di laurea magistrale in Computer Science and Engineering, 10 ore
- I semestre 2017/2018: Attività di didattica integrativa sperimentale (responsabile di laboratorio) per il corso Informatica B, AA 2017/2018, Politecnico di Milano, corso di laurea in Ingegneria Meccanica, 9 ore
- II semestre 2015/2016: Attività di tutorato (tutor di laboratorio) per il corso Prova Finale-Ingegneria del Software, AA 2015/2016, Politecnico di Milano, corso di laurea in Ingegneria Informatica, 32 ore

DOCUMENTATA ATTIVITÀ DI FORMAZIONE O DI RICERCA PRESSO QUALIFICATI ISTITUTI ITALIANI O STRANIERI;

- settembre 2023 - oggi: ricercatore presso Politecnico di Milano
- 8 giugno 2021 - 31 agosto 2023: ricercatore post dottorato presso Universitat Pompeu Fabra, Barcellona, Spagna
- 1 aprile - 31 maggio 2021: attività di supporto alla ricerca presso Politecnico di Milano
- 16 novembre 2017 - 11 marzo 2021: corso di dottorato in Ingegneria dell'Informazione/Information Technology, Politecnico di Milano
- 16 novembre 2017 - 11 marzo 2021: Assegnista di ricerca presso Politecnico di Milano
- 24 luglio - 2 agosto 2018: CIFAR Deep Learning and Reinforcement Learning Summer School, Toronto, Canada
- 7-14 ottobre 2017: ACAI Summer School on Reinforcement Learning, Nieuwpoort, Belgio
- 2015-2017: Corso di laurea magistrale in Computer Science and Engineering/Ingegneria Informatica, Politecnico di Milano

- 2012-2015: Corso di laurea in Ingegneria Informatica, Politecnico di Milano
- I semestre 2016/2017: Programma Erasmus, KTH Royal Institute of Technology, Stoccolma, Svezia (Computer Science)

REALIZZAZIONE DI ATTIVITÀ PROGETTUALE

- 2023 - oggi: partecipazione a *Artificial Intelligence Foundations for Sequential Decision Making*, Extended Partnership - Future Artificial Intelligence Research (FAIR). National Recovery and Resilience Plan, Mission 4 "Education and research" - Component 2 "From research to business" - Investment 1.3, funded by the European Union - NextGenerationEU. Presso Politecnico di Milano. PI: Nicola Gatti
- 2021: partecipazione a *Reinforcement learning techniques for developing artificial test drivers on a F1 simulator*, presso Politecnico di Milano, progetto industriale con Gestione Sportiva Ferrari (membro). PI: Marcello Restelli
- 2021-2022: partecipazione a *Provably Efficient Algorithms for Large-Scale Reinforcement Learning*, ERC Grant agreement No. 950180. Presso Universitat Pompeu Fabra, Barcellona, Spagna. PI: Gergely Neu
- 2017-2019: partecipazione a TOTAL EFFICIENCY 4.0 (POR FESR), presso Politecnico di Milano, progetto industriale con Pirelli Tyre S.p.A.. PI: Marcello Restelli

ORGANIZZAZIONE, DIREZIONE E COORDINAMENTO DI GRUPPI DI RICERCA NAZIONALI E INTERNAZIONALI, O PARTECIPAZIONE AGLI STESSI

- Dal 2017 al 2021 e dal 2023 ad oggi: partecipazione a gruppo di ricerca presso AIRLAB, Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano. Ruolo: ricercatore. PI: Marcello Restelli
- Dal 2021 al 2023: partecipazione a gruppo di ricerca presso Artificial Intelligence and Machine Learning Research Group, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcellona, Spagna. Ruolo: ricercatore. PI: Gergely Neu

ATTIVITÀ DI RELATORE A CONGRESSI E CONVEGNI NAZIONALI E INTERNAZIONALI

Su invito:

- 12 gennaio 2024 (invited speaker) "Mini-Workshop on Reinforcement Learning", University of Mannheim, Germany.
- 25 maggio 2023 (invited speaker) "Theory of Reinforcement Learning Workshop", University of Alberta, Edmonton, Canada
- 11 maggio 2023 (invited speaker) "AI Seminars", Politecnico di Milano
- 16 settembre 2022 (invited speaker) "Reinforcement Learning Conference, Technische Universität Dresden, Dresda, Germania
- 1 luglio 2021 (invited speaker) "Mathematical Statistics and Learning", Barcelona Graduate School of Economics, Barcellona, Spagna

- 19 settembre 2019 (invited speaker) “Workshop on Markets, Algorithms, Prediction and LEarning (MAPLE)”, Politecnico di Milano

Come autore in conferenze internazionali:

- Neurips 2021 (online edition): spotlight talk (registrato)
- ICML 2021 (online edition): oral (registrato)
- AISTATS 2020 (online edition): oral (registrato)
- ICML 2019: oral (Long Beach, USA)
- Neurips 2018: oral – top 3% (Montreal, Canada)
- ICML 2018: oral (Stoccolma, Svezia)

CONSEGUIMENTO DI PREMI E RICONOSCIMENTI NAZIONALI E INTERNAZIONALI PER ATTIVITÀ DI RICERCA

- 16 novembre 2022: Ayudas Juan de la Cierva-Formación 2021, borsa per ricercatori post-dottorato spagnola bandita dalla Agencia Estatal de Investigación e finanziata da NextGeneration EU. Importo: 64.800€ per due anni. Ruolo: ricercatore (Investigador). *Terzo in graduatoria nazionale per l'area “Information and communication Technologies”*
- Da febbraio 2022: membro di ELLIS (European Laboratory for Learning and Intelligent Systems)
- 2021: Neurips outstanding reviewer award (top 8%)
- 2020: ICML top reviewer certificate of appreciation (top 33%)
- 2019: ICML travel award (1300\$)
- 2018: oral a Neurips (top 3%)
- 2018: Neurips travel award (1000\$)
- 2018: ICML travel award (1500\$)
- 2017: Neurips travel award (1200\$)

TITOLI DI CUI ALL'ARTICOLO 24 COMMA 3 LETTERA A) E B) DELLA LEGGE 30 DICEMBRE 2010, N. 240

Da 1 settembre 2023: Ricercatore a tempo Determinato tipo A (RTDA) presso Politecnico di Milano. Fine contratto: 31/08/2026

N.B. Ai fini del presente bando dell'Università degli Studi di Milano, rientro tra i *candidati che per almeno trentasei mesi, anche cumulativamente, abbiano frequentato corsi di dottorato di ricerca o svolto attività di ricerca sulla base di formale attribuzione di incarichi, escluse le attività a titolo gratuito, presso università o istituti di ricerca, italiani o stranieri, diversi da quella che ha emanato il bando.*

PRODUZIONE SCIENTIFICA

PROFILI AUTORE

Google Scholar: h-index 13 <https://scholar.google.it/citations?user=A2WxZlsAAAAJ>

Scopus: 256 citazioni, 19 documenti, h-index 10, 25 co-autori

SOMMARIO PUBBLICAZIONI

Rivista/Conferenza	Numero pubblicazioni	Rating
JMLR	1	Scimago Q1, CORE A*
Machine Learning	2	Scimago Q1, CORE A
NeurIPS	5	CORE A*, GGS A++
ICML	5	CORE A*, GGS A++
AAAI	2	CORE A*, GGS A++
IJCAI	2	CORE A*, GGS A++
COLT	2	CORE A*, GGS A+
AISTATS	2	CORE A, GGS A+
IJCNN	1	CORE B, GGS A-
ALT	2	CORE B, GGS B

PUBBLICAZIONI SCIENTIFICHE

Journal Papers

[J1] G. Paczolay, M. Papini, A. M. Metelli, I. Harmati, and M. Restelli. Sample Complexity of Variance-Reduced Policy Gradient: Weaker Assumptions and Lower Bounds. *Machine Learning*, 2024 (to appear)

[J2] M. Papini, M. Pirotta, and M. Restelli. Smoothing policies and safe policy gradients. *Machine Learning*, 2022, 1-57 (<https://rdcu.be/c0a7S>)

[J3] A. M. Metelli, M. Papini, N. Montali, M. Restelli. Importance Sampling Techniques for Policy Optimization. *Journal of Machine Learning Research (JMLR)* 21.141., pp. 1-75, 2020

Conference Papers

[C1] M. Papini, G. Manganini, A. M. Metelli, and M. Restelli. Policy gradient with active importance sampling. In *Reinforcement Learning Conference (RLC)*, 2024 (to appear).

[C2] G. Neu, M. Papini, and L. Schwartz. Optimistic information directed sampling. In *COLT*, 2024 (to appear).

- [C3] A. Montenegro, M. Mussi, A. Metelli, and M. Papini. Learning optimal deterministic policies with stochastic policy gradients. In ICML, 2024 (to appear).
- [C4] D. Maran, A. Metelli, M. Papini, and M. Restelli. Projection by convolution: Optimal sample complexity for rl in continuous-space mdps. In COLT, 2024 (to appear).
- [C5] D. Maran, A. Metelli, M. Papini, and M. Restelli. No-regret reinforcement learning in smooth mdps. In ICML, 2024 (to appear).
- [C6] F. Bacchiocchi, F. E. Stradi, M. Papini, A. M. Metelli, and N. Gatti. Online learning with off-policy feedback in adversarial mdps. In IJCAI, 2024 (to appear).
- [C7] G. Gabbianelli, G. Neu, M. Papini, and N. Okolo. Offline primal-dual reinforcement learning for linear mdps. In AISTATS, volume 238 of Proceedings of Machine Learning Research, pages 3169–3177. PMLR, 2024.
- [C8] G. Gabbianelli, G. Neu, and M. Papini. Importance-weighted offline learning done right. In ALT, volume 237 of Proceedings of Machine Learning Research, pages 614–634. PMLR, 2024.
- [C9] G. Gabbianelli, G. Neu, and M. Papini. Online learning with off-policy feedback. In ALT, volume 201 of Proceedings of Machine Learning Research, pages 620–641. PMLR, 2023.
- [C10] A. Tirinzoni, M. Papini, A. Touati, A. Lazaric, and M. Pirotta. Scalable representation learning in linear contextual bandits with constant regret guarantees. Advances in Neural Information Processing Systems, 35:2307–2319, 2022.
- [C11] G. Neu, I. Olkhovskaia, M. Papini, and L. Schwartz. Lifting the information ratio: An information-theoretic analysis of thompson sampling for contextual bandits. Advances in Neural Information Processing Systems, 35:9486–9498, 2022.
- [C12] M. Papini, A. Tirinzoni, M. Restelli, A. Lazaric, and M. Pirotta. Leveraging good representations in linear contextual bandits. In ICML, volume 139 of Proceedings of Machine Learning Research, pages 8371–8380. PMLR, 2021.
- [C13] M. Papini, A. Tirinzoni, A. Pacchiano, M. Restelli, A. Lazaric, and M. Pirotta. Reinforcement learning in linear mdps: Constant regret and representation selection. In NeurIPS, pages 16371–16383, 2021.
- [C14] A. M. Metelli, M. Papini, P. D’Oro, and M. Restelli. Policy optimization as online learning with mediator feedback. In AAAI, pages 8958–8966. AAAI Press, 2021.
- [C15] M. Papini, A. Battistello, and M. Restelli. Balancing learning speed and stability in policy gradient via adaptive exploration. In AISTATS, volume 108 of Proceedings of Machine Learning Research, pages 1188–1199. PMLR, 2020.
- [C16] P. D’Oro, A. M. Metelli, A. Tirinzoni, M. Papini, and M. Restelli. Gradient-aware model-based policy search. In AAAI, pages 3801–3808. AAAI Press, 2020.
- [C17] L. Bisi, L. Sabbioni, E. Vittori, M. Papini, and M. Restelli. Risk-averse trust region optimization for reward-volatility reduction. In IJCAI, pages 4583–4589. ijcai.org, 2020.
- [C18] M. Papini, A. M. Metelli, L. Lupo, and M. Restelli. Optimistic policy optimization via multiple importance sampling. In ICML, volume 97 of Proceedings of Machine Learning Research, pages 4989–4999. PMLR, 2019.
- [C19] M. Beraha, A. M. Metelli, M. Papini, A. Tirinzoni, and M. Restelli. Feature selection via mutual information: New theoretical insights. In IJCNN, pages 1–9. IEEE, 2019.

[C20] M. Papini, D. Binaghi, G. Canonaco, M. Pirotta, and M. Restelli. Stochastic variance-reduced policy gradient. In ICML, volume 80 of Proceedings of Machine Learning Research, pages 4023–4032. PMLR, 2018.

[C21] A. M. Metelli, M. Papini, F. Faccio, and M. Restelli. Policy optimization via importance sampling. In NeurIPS, pages 5447–5459, 2018.

[C22] M. Papini, M. Pirotta, and M. Restelli. Adaptive batch size for safe policy gradients. In NeurIPS, pages 3591–3600, 2017.

Workshop Papers

[W1] G. Neu, M. Papini, L. Schwartz. Optimistic Information Directed Sampling. FoRLaC (Foundations of Reinforcement Learning and Control) workshop at ICML, Vienna, Austria, 2024.

[W3] F. Bacchiocchi, F. E. Stradi, M. Papini, A. M. Metelli, N. Gatti. Online Adversarial MDPs with Off-Policy Feedback and Known Transitions. 16th European Workshop on Reinforcement Learning, Brussels, Belgium, 2023

[W4] G. Gabbianelli, G. Neu, N. Okolo, M. Papini. Offline Primal-Dual Reinforcement Learning for Linear MDPs. 16th European Workshop on Reinforcement Learning, Brussels, Belgium, 2023

[W5] G. Neu, J. Olkhovskaya, M. Papini and L. Schwartz. Lifting the Information Ratio: An Information-Theoretic Analysis of Thompson Sampling for Contextual Bandits. 15th European Workshop on Reinforcement Learning, Milan, Italy, 2022

[W6] A. Tirinzoni, M. Papini, A. Touati, A. Lazaric, and M. Pirotta. Scalable Representation Learning in Linear Contextual Bandits with Constant Regret Guarantees. 15th European Workshop on Reinforcement Learning, Milan, Italy, 2022

[W7] G. Gabbianelli, M. Papini, G. Neu. Online Learning with Off-Policy Feedback. ICML-2022 workshop on Complex Feedback in Online Learning, Baltimore, USA, 2022

[W8] A. Gianola, M. Montali, and M. Papini. Automated Reasoning for Reinforcement Learning Agents in Structured Environments. OVERLAY workshop on fOrmal VERification, Logic, Automata and sYNthesis, Padova, Italy, 2021

[W9] M. Papini, A. Tirinzoni, A. Pacchiano, M. Restelli, A. Lazaric, and M. Pirotta. Reinforcement Learning in Linear MDPs: Constant Regret and Representation Selection. ICML Workshop on Reinforcement Learning Theory, virtual, 2021

[W10] M. Papini, A. Battistello, and M. Restelli. Safe Exploration in Gaussian Policy Gradient. NeurIPS-2019 Workshop on Safety and Robustness in Decision Making, Vancouver, Canada, 2019

[W11] M. Papini, A. Battistello, and M. Restelli. Safely Exploring Policy Gradient. 14th European Workshop on Reinforcement Learning, Lille, France, 2018

LAVORI IN PREPARAZIONE

Articoli in Peer-Review

- A. Montenegro, M. Mussi, M. Papini, A. M. Metelli. Last-Iterate Global Convergence of Policy Gradients for Constrained Reinforcement Learning. 2024.
- D. Maran, A. M. Metelli, M. Papini, M. Restelli. Local Linearity: the Key for No-regret Reinforcement Learning in Continuous MDPs. 2024.
- G. Tedeschi, M. Papini, A. M. Metelli, M. Restelli. Search or Split: Policy Gradient with Adaptive Policy Spaces. 2024.
- G. Paczolay, M. Papini, A. M. Metelli, I. Harmati, M. Restelli. Stabilizing Policy Gradient with Active Importance Sampling. 2024.

Collaborazioni Attive

- Con Nicolò Cesa-Bianchi, Università degli Studi di Milano, su *Reinforcement Learning with Human Feedback*
- Con Alessandro Gianola, Técnico Lisboa, su *Symbolic Reinforcement Learning*
- Con Laura Toni, University College London, su *Graph-based Exploration in Reinforcement Learning*
- Con Dimitri Ognibene, Università degli Studi di Milano-Bicocca, su *Reinforcement Learning and Addiction*

ALTRO

ORGANIZZAZIONE DI EVENTI SCIENTIFICI

- 2025: **local chair** di International Conference on Algorithmic Learning Theory (ALT) presso Politecnico di Milano (to be announced)
- 2023: organizzatore (main organizer con Vincent Adam) di Reinforcement Learning Summer School (RLSS), 28 giugno-5 luglio 2023, Universitat Pompeu Fabra, Barcellona, Spagna
- 2023: co-organizzatore di ELLIS Pre-NeurIPS Fest, 4 dicembre 2023, ELLIS Unit of Milan

ATTIVITA' EDITORIALE

- 2024: **Area Chair** per NeurIPS
- 2024: Reviewer per ICML, COLT, Senior Reviewer per RLC (Reinforcement Learning Conference)
- 2022-2024: Reviewer for Transactions on Machine Learning Research (TMLR)
- 2023: Reviewer per ICML, COLT, EWRL, NeurIPS, ICLR, AISTATS
- 2022: Reviewer per IEEE Transactions on Automatic Control (TACON)
- 2022: Reviewer per ICML, NeurIPS, EWRL
- 2021 Outstanding Reviewer (top 8%) per NeurIPS, expert reviewer per ICML, reviewer per AISTATS, emergency reviewer per AAAI

- 2020: Top 33% reviewer per ICML, reviewer per NeurIPS, AISTATS, AAAI, UAI, ECAI
- 2019: Reviewer per ICML, NeurIPS, UAI

PARTECIPAZIONE A COMITATI

- 27 giugno 2024: membro commissione di difesa di dottorato, PhD program in Information Technology, Politecnico di Milano (nominato)
- maggio 2024: membro commissione bando a cascata del progetto "Future Artificial Intelligence Research -- FAIR", (Bando Imprese), codice PE000013, PNRR, Missione 4, Componente 2, Investimento 1.3 (nominato)
- 19 dicembre 2023: membro commissione lauree magistrali, Computer Science and Engineering, Politecnico di Milano
- 30 marzo 2023: membro commissione Thesis Proposal Defenses, PhD program in Information and Communication Technologies, Universitat Pompeu Fabra, Barcellona, Spagna
- 2023: membro della giuria per il premio CLAIRE R2Net 2022 Papers Highlights
- 2021, 2022: "Evaluator" per la selezione delle applicazioni a ELLIS PhD program
- July 5 2022: membro di commissione lauree , Universitat Pompeu Fabra, Barcelona, Spain
- 17 marzo 2022: membro commissione Thesis Proposal Defense, PhD program in Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

SUPERVISIONE DI STUDENTI

- 2023 - oggi: co-supervisione di un dottorando (Alessandro Montenegro), Information Technology, Politecnico di Milano. Co-supervisore: Alberto Maria Metelli
- 2023 - oggi: supervisione di 4 studenti di laurea magistrale e co-supervisione di altri 3, Computer Science and Engineering, Politecnico di Milano
- 2023 - oggi: mentoring di 5+ dottorandi presso Politecnico di Milano
- 2021 - 2023: mentoring di 3 dottorandi presso Universitat Pompeu Fabra
- 2018 - 2020: correlatore di 8 tesi di laurea magistrale, Politecnico di Milano

PARTECIPAZIONE A CONFERENZE E WORKSHOP

- 2024: ICML, Vienna, Austria, 2 poster (planned)
- 2024: AISTATS, Valencia, Spagna, 1 poster
- 2024: ALT, San Diego, USA, 1 poster
- 2024: Mini Workshop on Reinforcement Learning, Mannheim, Germania, 1 invited talk
- 2023: EWRL 16, Brussels, Belgio, 2 poster
- 2023: AISTATS, Valencia, Spagna

- 2023: Upper Bound, Edmonton, Canada, invited talk a RL Theory Workshop
- 2023: ALT, Singapore, come co-autore e session chair
- 2022: NeurIPS, New Orleans, USA 2 poster
- 2022: ELLIS ILIR Workshop, Feldberg, Germania
- 2022: EWRL 15, Milano, 2 poster
- 2022: Reinforcement Learning Conference, Dresda, Germania, invited talk
- 2022: ICML-2022 Workshop on Complex Feedback in Online Learning, Baltimore, USA, 1 poster
- 2021: NeurIPS (online edition) 1 poster (live online) e 1 spotlight talk (registrato)
- 2021: Mathematical Statistics and Learning, Barcellona, Spagna (invited talk)
- 2021: ICML (online edition), 1 poster (live online) e 1 oral (registrato)
- 2021: AAAI (online edition), 1 poster (live online)
- 2021: IJCAI (online edition), 1 poster (live online)
- 2020: NeurIPS (online edition)
- 2020: AISTATS (online edition), 1 oral (registrato)
- 2019: Workshop MAPLE, Milano, invited talk
- 2019: ICML, Long Beach, USA, 1 poster e 1 oral
- 2018: NeurIPS, Montreal, Canada, 1 poster e 1 oral
- 2018: EWRL 14, Lille, France, 1 poster
- 2018: ICML, Stoccolma, Svezia, 1 poster, 1 oral
- 2017: NeurIPS, Long Beach, USA, 1 poster

RESEARCH STATEMENT

ON-GOING RESEARCH AND RECENT ACHIEVEMENTS

Research focus. My research is focused on **reinforcement learning theory and algorithms**. Reinforcement Learning (RL) [17] is a framework for *artificial intelligence* and an application of *machine learning* techniques to the problem of sequential *decision making* under uncertainty, which encompasses a wide range of decision-making and *control* applications: from robotics to online advertising, from videogames to personalized medicine. All of these problems can, in principle, be modeled as the ongoing interaction of an adaptive agent with an unknown environment. Reinforcement learning algorithms are general methods to equip artificial agents with the ability to improve from direct experience. The theory of RL is concerned with the efficiency of the learning process and with the performance, generalization capabilities, and safety of the behavior *policy* learned by the agent.

Distinctive aspects of my research. I have worked on reinforcement learning since my PhD, and before. I have worked on many aspects of RL, but my research is and has always been characterized by the following key aspects:

- A focus on *large-scale* problems, that is, decision problems with many states and actions, or even **continuous control** problems, that cannot be described by discrete models.
- A characterization of RL problems in terms of **sample complexity**, that is, the amount of data (interaction episodes) that is required to find a good solution.
- The combination of canonical RL techniques with methods from **optimization**, especially of the stochastic non-convex kind, and **online learning**, especially *contextual bandit algorithms* [18].

Summary of main results (2017-2021). The main achievements of my research activity can be summarized as follows and divided into the following areas. This is work that I have done during my PhD at Politecnico di Milano and at Facebook AI Research (now Meta) with several co-authors.

- **Sample efficiency:** I have studied the sample complexity of *policy gradient methods*, a family of RL algorithms particularly suited for continuous control problems. I have developed the first such method with an improved sample complexity [1] over the standard algorithm [19], adapting relatively new *variance-reduction* techniques from the *stochastic optimization* literature [20]. In doing so, I laid down the foundations for further developments, essentially opening a line of research on variance-reduced policy gradient algorithms [e.g., 21,22,23]. I have also studied ways to employ *importance-sampling* estimation to improve the sample efficiency of policy optimization algorithms [2,6,8,12].
- **Safety:** I have studied the *monotonic improvement* properties of policy gradient methods. Monotonic improvement guarantees that the performance of the agent does not oscillate during the learning process, preventing potentially dangerous deviations from safe behavior. I have designed policy gradient algorithms with provable monotonic improvement guarantees [4, 5]. I have also studied the related problems of *safe exploration* [13] and *risk-averse* reinforcement learning [11].
- **Representation learning:** I have provided a formal characterization of good *representations* (state-action features) for reinforcement learning, in the sense of representations that allow sample-efficient learning. I have developed methods to select good representations in an online fashion. Starting from the relatively simple setting of linear contextual bandits [3], I have generalized my findings to the full reinforcement learning problem [9] and to large-scale problems using *deep neural networks* as feature extractors [10].

Recent achievements (since 2021). My recent research is characterized by two main aspects of novelty compared to my older works, listed in the following. This change reflects my period as a postdoctoral researcher at Universitat Pompeu Fabra, Barcelona, in a research group that is even more focused on the theoretical aspects of reinforcement learning. This is work that I have done there with international co-authors.

- I have dedicated more time to the study of **contextual bandit** problems. This can be seen as a simplified version of RL where the choices of the agent do not have long-term effects. Targeted online advertising is the typical application example. However, I see contextual bandits it mainly as a testing ground to study some fundamental aspects of reinforcement learning, such as the *exploration-exploitation dilemma* and the role of representation learning and *function approximation* in large-scale problems.
I have contributed directly to the theory of contextual bandits by generalizing a famous information-theoretic analysis of Thompson sampling [24], an extremely popular algorithm for multi-armed bandits [18], to the challenging setting of contextual bandits with nonlinear rewards and adversarial contexts [7].
- My previous research was entirely on *online* RL, where the agent interacts with the environment in real time and must actively *explore* its available choices. Recently, I have started working also on **offline reinforcement learning** [25]. In this scenario, online trial-and-error is replaced by access to a *dataset* of previous interactions. I have studied this problem both for contextual bandits [14,16] and large-scale RL [15].

Current research (since 2023). I am currently focusing on applying the theoretical skills that I have recently acquired to a favorite problem of mine, the sample complexity of policy gradient methods. At the same time, I am continuing to work on the theory of offline RL, which is a very active area of research, and resuming my work on safe RL and representation learning.

Finally, motivated by the recent astounding advancements in the development of chatbots [26], I have started exploring the problem of *RL by human feedback (RLHF)*. Here is a list of active projects that I am conducting with co-authors and (co-)supervised students from Politecnico di Milano, and co-authors from other institutions, organized into the previously mentioned areas.

- **Sample efficiency**
 - *Advancements in variance-reduced policy gradient*: with the ambitious goal of closing the line of research that I have started back in 2018 [1], we establish optimal sample-complexity results under weaker assumptions and with matching lower-bounds (just accepted by Machine Learning journal).
 - *Active importance sampling for policy gradient methods*: we study an alternative way to improve sample efficiency by designing good exploration policies (one paper accepted at Reinforcement Learning Conference, one submitted to NeurIPS, targeting AAAI next).
 - *Parameter based exploration*: we study an alternative policy optimization method that is widespread in robotic applications, but has been mostly ignored by theoreticians. (ongoing).
- **Safety**:
 - *Constrained RL*: we study the problem of learning under constraints in continuous control problems. (accepted to ICML 2024).
 - *Deterministic RL*: we characterize the situations in which random exploration, which is inherently unsafe, can be avoided while preserving sample efficiency. (ongoing).
- **Representation learning**: we study alternative representations for continuous control problems based on orthogonal polynomials (two papers accepted at ICML 2024 and COLT 2024, one submitted to NeurIPS).
- **Contextual bandits...with human feedback**: we study an online learning problem human-like feedback (ongoing).
- **Offline RL**: we are tackling the long-standing problem of establishing meaningful worst-case lower bounds (ongoing)

SHORT-TO-MID-TERM RESEARCH DIRECTIONS

Vision. The field of artificial intelligence and machine learning has seen a fast, tremendous development in the last decade, powered by the availability of huge amounts of data and computational resources that allowed deep neural networks to fully express their potential. The scene was dominated first by image-processing applications, and more recently by *large language models* [26]. Besides some

striking results on games [27], reinforcement learning has struggled so far to find widespread application, in spite of its promised generality. I believe that the following **five years** will be crucial to advance reinforcement learning from a mere research subject to a transformative technology. I also believe that reinforcement learning *theory* will play a fundamental role in the design of powerful yet reliable algorithms. However, theory will have to confront itself with the concrete needs and methodologies of real-world applications.

Directions. I identify the following research directions at the intersection of my theoretical expertise and the contemporary challenges of artificial intelligence.

- Data for reinforcement learning, that is, interaction data, are harder to obtain compared to supervised or unsupervised learning applications. Hence, to keep up with the incoming challenges of artificial intelligence, the sample efficiency of RL algorithms must be pushed further. Thus, the theory of **sample complexity** is more relevant than ever.
- Future RL algorithms will need to process data coming from several different sources: direct experience, simulation, historical data, *human feedback*. Tools from the field of **offline reinforcement learning** can be used to combine these diverse data in the most efficient manner.
- While large language models seem capable of solving a wide range of tasks in the virtual world, the interface with the real world remains one of the greatest challenges of artificial intelligence. I think that the greatest potential of reinforcement learning lies precisely in cyberphysical systems, such as robots and autonomous vehicles. From the theoretical perspective, a shift to **continuous control** problems is absolutely necessary.
- Contemporary control and decision system must be able to process high-dimensional input data like images and language. **Representation learning** is fundamental to extract the relevant information from these data and make the best decisions based on them.
- As artificial intelligence systems become more and more present in our industry and day-to-day life, the **safety** of these systems will assume a fundamental importance. I am particularly concerned with the safety of adaptive systems that keep learning after they have been deployed. Combining online learning and safety is a great theoretical and technological challenge.

Projects. All of the aforementioned research directions can be developed into concrete projects within a five years span. Here I present two projects that are already well defined and are likely to produce significant results on a shorter term:

- **Hybrid online/offline RL:** *offline reinforcement learning* allows to exploit large amounts data of past interactions, like logs of the activity of the very legacy systems that we aim to replace. However, it does not consider any direct interaction between the learning agent and the environment: once the historical data have been processed, the learning is over. The nascent field of hybrid RL [28] allows to combine offline data and online interaction. Going further, we can use this framework to design *lifelong learning* agents that are capable of integrating data from different sources: starting from the background knowledge represented by the historical data, use the information obtained so far to guide the behavior of the agent *and* the collection of more data. The theory of offline RL offers the tools to characterize the importance and quality of data, but we need to integrate it with the theory of *exploration* in online reinforcement learning. Moreover, novel algorithmic ideas are necessary for the hybrid setting.
I plan to hire a PhD student in the near future to work on this project, and a master student is already working on some aspects of it.
I have also applied to the FIS Starting Grant proposing a project on hybrid RL.
- **Theory of parameter-based policy optimization:** parameter-based policy optimization [29] is one of the most used reinforcement algorithms in *continuous control* applications, such as industrial robotics. The reasons can be found in its simplicity, robustness to the noise and the nonstationarity of the environment, and the ease of integrating domain knowledge into the learning process. However, this family of algorithms has been mostly ignored by the RL theory

community so far. Besides their simplicity of implementation, a rigorous characterization of the convergence, sample complexity, safety and scalability of this algorithms is highly non-trivial, and will likely require less common tools and novel ideas.

One PhD student is starting to work on this. This is a joint effort with other members of the RL research group.

Methodology. I will address the aforementioned problems with the same combination of theory and algorithmic design that led me to the achievements mentioned in the previous section.

While the motivation typically comes from real-world applications, the first step is always to model the problem in a clear and rigorous manner, trying to make only those simplifying assumptions that are strictly necessary. The **theoretical analysis** guides the design of the algorithm, in what is often an iterative creative process. Sometimes, this consists in the refinement of an existing algorithm. Once the final algorithm is determined, the theory is refined to formally establish its properties of convergence, sample complexity, or safety, typically in terms of probabilistic guarantees.

In most cases, the end goal is the actual implementation of a prototype **algorithm** that can be tested on benchmark problems or on the motivating application. In this phase, some compromises are often necessary with respect to the theory, due to simplifying assumptions that fail to hold in the real world or reasons of *computational* efficiency and scalability. Finally, **experiments** are designed to empirically measure the properties of the concrete algorithm, its performance, and to acknowledge the remaining gap between theory and practice. Results are reported together with their statistical significance and all the information that is necessary to replicate them.

REFERENCES

- [1] Papini, M., Binaghi, D., Canonaco, G., Pirotta, M., & Restelli, M. (2018, July). Stochastic variance-reduced policy gradient. In International conference on machine learning (pp. 4026-4035). PMLR.
- [2] Metelli, A. M., Papini, M., Faccio, F., & Restelli, M. (2018). Policy optimization via importance sampling. *Advances in Neural Information Processing Systems*, 31.
- [3] Papini, M., Tirinzoni, A., Restelli, M., Lazaric, A., & Pirotta, M. (2021, July). Leveraging good representations in linear contextual bandits. In International Conference on Machine Learning (pp. 8371-8380). PMLR.
- [4] Papini, M., Pirotta, M., & Restelli, M. (2017). Adaptive batch size for safe policy gradients. *Advances in neural information processing systems*, 30.
- [5] Papini, M., Pirotta, M., & Restelli, M. (2022). Smoothing policies and safe policy gradients. *Machine Learning*, 111(11), 4081-4137.
- [6] Metelli, A. M., Papini, M., Montali, N., & Restelli, M. (2020). Importance sampling techniques for policy optimization. *The Journal of Machine Learning Research*, 21(1), 5552-5626.
- [7] Neu, G., Olkhovskaia, I., Papini, M., & Schwartz, L. (2022). Lifting the information ratio: An information-theoretic analysis of thompson sampling for contextual bandits. *Advances in Neural Information Processing Systems*, 35, 9486-9498.
- [8] Papini, M., Metelli, A. M., Lupo, L., & Restelli, M. (2019, May). Optimistic policy optimization via multiple importance sampling. In International Conference on Machine Learning (pp. 4989-4999). PMLR.
- [9] Papini, M., Tirinzoni, A., Pacchiano, A., Restelli, M., Lazaric, A., & Pirotta, M. (2021). Reinforcement learning in linear mdps: Constant regret and representation selection. *Advances in Neural Information Processing Systems*, 34, 16371-16383.
- [10] Tirinzoni, A., Papini, M., Touati, A., Lazaric, A., & Pirotta, M. (2022). Scalable representation learning in linear contextual bandits with constant regret guarantees. *Advances in Neural Information Processing Systems*, 35, 2307-2319.
- [11] Bisi, L., Sabbioni, L., Vittori, E., Papini, M., & Restelli, M. (2019). Risk-averse trust region optimization for reward-volatility reduction. arXiv preprint arXiv:1912.03193.
- [12] Metelli, A. M., Papini, M., D'Oro, P., & Restelli, M. (2021, May). Policy optimization as online learning with mediator feedback. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 35, No. 10, pp. 8958-8966).

- [13] Papini, M., Battistello, A., & Restelli, M. (2020, June). Balancing learning speed and stability in policy gradient via adaptive exploration. In International conference on artificial intelligence and statistics (pp. 1188-1199). PMLR.
- [14] Gabbianelli, G., Neu, G., & Papini, M. (2023, February). Online learning with off-policy feedback. In International Conference on Algorithmic Learning Theory (pp. 620-641). PMLR.
- [15] Gabbianelli, G., Neu, G., Okolo, N., & Papini, M. (2023). Offline Primal-Dual Reinforcement Learning for Linear MDPs. arXiv preprint arXiv:2305.12944.
- [16] Gabbianelli, G., Neu, G., & Papini, M. (2024). Importance-Weighted Offline Learning Done Right. ALT 2024.
- [17] Sutton, Richard S., and Andrew G. Barto. Reinforcement learning: An introduction. MIT press, 2018.
- [18] Lattimore, Tor, and Csaba Szepesvári. Bandit algorithms. Cambridge University Press, 2020.
- [19] Williams, Ronald J. "Simple statistical gradient-following algorithms for connectionist reinforcement learning." Machine learning 8 (1992): 229-256.
- [20] Johnson, Rie, and Tong Zhang. "Accelerating stochastic gradient descent using predictive variance reduction." Advances in neural information processing systems 26 (2013).
- [21] Xu, Pan, Felicia Gao, and Quanquan Gu. "Sample Efficient Policy Gradient Methods with Recursive Variance Reduction." International Conference on Learning Representations. 2019.
- [22] Zhang, Junyu, et al. "On the convergence and sample efficiency of variance-reduced policy gradient method." Advances in Neural Information Processing Systems 34 (2021): 2228-2240.
- [23] Gargiani, Matilde, et al. "PAGE-PG: A simple and loopless variance-reduced policy gradient method with probabilistic gradient estimation." International Conference on Machine Learning. PMLR, 2022.
- [24] Russo, D., & Van Roy, B. (2016). An information-theoretic analysis of thompson sampling. The Journal of Machine Learning Research, 17(1), 2442-2471.
- [25] Levine, S., Kumar, A., Tucker, G., & Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems. arXiv preprint arXiv:2005.01643.
- [26] OpenAI (2023). GPT-4 Technical Report. ArXiv, abs/2303.08774.
- [27] Wurman, P. R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T. J., ... & Kitano, H. (2022). Outracing champion Gran Turismo drivers with deep reinforcement learning. Nature, 602(7896), 223-228.
- [28] Song, Y., Zhou, Y., Sekhari, A., Bagnell, J. A., Krishnamurthy, A., & Sun, W. (2022). Hybrid rl: Using both offline and online data can make rl efficient. arXiv preprint arXiv:2210.06718.
- [29] Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., & Schmidhuber, J. (2010). Parameter-exploring policy gradients. Neural Networks, 23(4), 551-559.

TEACHING STATEMENT

SHORT-TO-MID-TERM TEACHING PLAN

It is my intention to contribute to the educational offer of this University by teaching classes on **machine learning** and **artificial intelligence** and introductory computer science classes.

Topics. My expertise is on **reinforcement learning**, so I would be particularly suited to teach advanced classes in this area. For instance, I could teach a PhD level class on some advanced topics such as: policy gradient algorithms, reinforcement learning theory, contextual bandit algorithms, Markov decision processes, or deep reinforcement learning.

I could as well teach classes on machine learning and artificial intelligence in general, at master or PhD level. Besides the canonical *machine learning*, *artificial intelligence*, *deep learning*, and *reinforcement learning* classes, an example of a novel theory-centered class would be **computational learning theory**. On the more practical side, I could teach a class on **automatic differentiation** with a focus on frameworks such as *tensorflow*, *pytorch* and *jax*. I could also host project-based classes on machine learning applications, or flipped-classroom courses where the students present machine learning papers.

Finally, I could teach **introductory computer science** classes at bachelor level, for instance on programming, foundations of computer science, algorithms.

Prior experience. During my PhD (2017-2021) and in my current position of Assistant Professor (since September 2023), I have worked as a **teaching assistant** at Politecnico di Milano for several classes. These include *introductory computer science*, *artificial intelligence*, and *machine learning* classes. My teaching experience starts even earlier, since I worked as a lab tutor during my master for the final project of the Computer Engineering bachelor.

I also volunteered as a teaching assistant for two international **Reinforcement Learning Summer Schools** (Lille 2019, Amsterdam 2022) targeted to master and PhD students.

Finally, I *organized* the last edition of the Reinforcement Learning Summer School (Barcelona 2023) while I was a postdoctoral researcher at Universitat Pompeu Fabra, Barcelona.

Next year (2024/2025) I will teach *my own class* on Algorithms in the Computer Engineering bachelor of Politecnico di Milano.

Principles. Thanks to my prior experience as a student, a teaching assistant, as a summer school organizer, I have matured several ideas on how teaching can be made effective and engaging that I can summarize in the following principles:

- **Empathy:** I believe that the main pitfall of teaching, especially at the university level, is the cognitive bias known as the *curse of knowledge*. When we are confident in our topics, it is easy to forget how it felt like to have no prior knowledge on the matter. As a teacher, I always try to see the subject from the perspective of the student.
- **Active thinking:** I think that students learn the most when they have to *think* about the concepts and methods that are presented, especially in engineering classes. I try to design the lecture in a way that encourages this process. The main way to achieve this is to show the thought process that leads to a solution *before* presenting the solution itself. I also try to avoid *slide presentations* whenever possible, using the **blackboard/tablet** for theory and **live-coding** for practice, to display the thought process in real time.
- **Interaction:** I think that a small but well targeted amount of interaction can greatly benefit the efficacy and the engagement of the lectures. I particularly rely on *brain-stormin*, letting the

students propose their solutions and critically analyzing them on the fly, even if they are obviously correct or incorrect.

- **Preparation/Improvisation:** I think that a good lecture must be prepared in advance, but some element of improvisation is also important, to deliver a faithful demonstration of the *thought processes* and to avoid the *curse of knowledge*, as mentioned in previous points. I try to prepare the material and the exercises with great care, but leave the way I will present them more open.
- **Engagement:** In preparing the lectures, I try not to forget the most practical aspects of engagement. A well placed short break, example, exercise, brain-storming session (even jokes and anecdotes, if used with parsimony) can have amazing effects on the level of attention of the students. A technique that I particularly like is giving the students a short break to *think* about a question. Another helpful technique is to alternate content of different kinds, such as theory and exercises, or intense blackboard writing and live coding.
- **Honest confidence:** I think that a teacher should be confident in the topics of his class. If the teacher seems confused, the students may feel completely lost and frustrated. Unfortunately, I acknowledge that it can take *years* of teaching the same subject to achieve true confidence, including confidently answering to the students' questions. As a teacher with less experience, I try *not to fake confidence*. If I am not sure about how to answer a question or I suddenly have some doubts about what I have prepared, I prefer to be honest about it, check later, and come back the next time with an informed answer or correction.
- **Positive attitude:** I think that the attitude of the teacher can significantly influence the mood of the students, and in turn their mood has an impact on their receptivity, and so ultimately on the efficacy of the whole lecture. For this reason I try to create a respectful but chill environment where the students can feel at ease as long as they respect some basic rules, and are not afraid to interact with the teacher. For the same reason I try to remedy to the students' knowledge gaps without making them feel guilty, encourage answers that may sound trivial or wrong, and questions that may sound naive. Finally, I encourage the students to face the problems with a realistic but positive attitude: the problems *are* hard, but if they fully take advantage of the lecture they *will* be able to solve them.
- **Fair exams:** I think that a good exam should be challenging but not punitive. That is, questions and exercises should be designed in a way that the students must really *apply* the concepts and methods learned during the class, and fail the exam when they fail to do so. However, questions should not be made difficult in an artificial way, for instance by adding confusing or ambiguous answers to multiple-choice questions, inserting easy-to-miss important details, making the exam longer, or requiring a disproportionate amount of brainless computations.

Method. For a **bachelor** level class, I would propose a balanced amount of theory and guided examples, presented on the blackboard/tablet or in live coding. Moreover, I would reserve a significant amount of hours to exercise and lab sessions. For applied classes, I would have a written exam and a practical one. For more theoretical classes, I would have a written or an oral exam. In both cases, I would propose an optional small project for extra points, unless the class is already project-based. For a **master** level class I would follow a similar scheme, but dedicate more time to interaction, individual exercises, and *group* projects.

I would organize a **PhD** level class either as a series of short seminars followed by a discussion, or as a flipped-classroom experience where the students present state-of-the-art topics. In this case I would have an oral evaluation for theory-centered classes and just a project for more practical ones, or use continual evaluation in the case of the flipped-classroom. I am open to experiment with *innovative teaching* even at the bachelor and master level, depending on the school's guidelines, the students preferences, and the logistic constraints. This may include flipped-classroom and continual evaluation.

LINGUE

Italiano: madrelingua

Inglese: C1 (Cambridge FCE grade A)

Spagnolo: intermedio

Data

20 giugno 2024

Luogo

Milano