

Riccardo De Santi

✉ rdesanti@ethz.ch, rdesanti@caltech.edu

🌐 riccardodesanti.com

📄 riccardo-de-santi

👤 Google Scholar

My research centers on **exploration for out-of-distribution discovery**: from foundational work on exploration in reinforcement learning to current work on **flow- and diffusion-based generative models** for discovering **new-to-nature chemical and biological structures**. More broadly, I aim to develop the principles that underlie discovery systems at the intersection of generative modeling, optimization, and decision-making under uncertainty, and ultimately contribute to the development of a **science of generative discovery**.

Selected Highlights

- July 2022 **Outstanding Paper Award at ICML 2022** on algorithmic foundations of exploration.
- December 2025 **Spotlight at NeurIPS 2025** on large-scale flow-based exploration for discovery.
- March 2023 **Awarded ETH AI Center Doctoral Fellowship** (~ 13/650 acceptance rate).
- March 2026 – present **Visiting Nobel Laureate Frances H. Arnold's and Yisong Yue labs at Caltech**.
- October 2023 **Given TedX talk "Reimagining scientific discovery with AI"** during MSc studies.

Selected Publications

- [1] **Riccardo De Santi**, Malte Franke, Ya-Ping Hsieh, and Andreas Krause. "A Unified Density Operator View of Flow Control and Merging". In: *International Conference on Machine Learning (ICML)*, **Oral presentation** at *Workshop on Real-World Constrained and Preference-Aligned Flow and Diffusion-Based Models at ICLR 2026*. (2026).
- [2] Sven Gutjahr^{*}, †, **Riccardo De Santi**^{*}, Luca Schaufelberger^{*}, Kjell Jorner, and Andreas Krause. "Constrained Molecular Generation via Sequential Flow Model Fine-Tuning". In: *International Conference on Machine Learning (ICML) 2026*, **Oral presentation** at *Workshop on Frontiers in Probabilistic Inference: Sampling meets Learning at NeurIPS*. 2026.
- [3] **Riccardo De Santi**^{*}, Kimon Protopoulos^{*}, Ya-Ping Hsieh, and Andreas Krause. "Verifier-Constrained Flow Expansion for Discovery Beyond the Data". In: *International Conference on Learning Representations (ICLR)*. 2026.
- [4] **Riccardo De Santi**, Marin Vlastelica, Ya-Ping Hsieh, Zebang Shen, Niao He, and Andreas Krause. "Flow Density Control: Generative Optimization Beyond Entropy-Regularized Fine-Tuning". In: **Oral presentation** at *Workshop on Generative AI and Biology at ICML 2025*. **Spotlight** at *Conference on Neural Information Processing Systems (NeurIPS)*. 2025.
- [5] **Riccardo De Santi**^{*}, Marin Vlastelica^{*}, Ya-Ping Hsieh, Zebang Shen, Niao He, and Andreas Krause. "Provable maximum entropy manifold exploration via diffusion models". In: *International Conference on Machine Learning (ICML)*. 2025.
- [6] Mirco Mutti^{*}, **Riccardo De Santi**^{*}, and Marcello Restelli. "The Importance of Non-Markovianity in Maximum State Entropy Exploration". In: *International Conference on Machine Learning (ICML)* **Outstanding Paper Award**. 2022.

^{*} denotes shared first-authorship, † denotes B.S. or M.S. student supervision.

Education

- Dec 2023 – present **Ph.D. Student**, *ETH AI Center, ETH Zurich*.
Recipient of the ETH AI Center Doctoral Fellowship. Working on diffusion- and flow-based exploration for out-of-distribution discovery: from theory, to scalable methods, and application on molecular and biological structures. Advised by Andreas Krause and co-advised by Niao He and Kjell Jorner.
- Sep 2020 – Nov 2023 **M.Sc. in Computer Science**, *ETH Zurich*.
Major in Machine Intelligence; minor in Theoretical Computer Science. During the M.Sc., published five papers on exploration in reinforcement learning at ICML, NeurIPS, ICLR, AAAI, and JMLR, including work recognized with the Outstanding Paper Award at ICML 2022.

Sep 2017 – Jul 2020 **B.Sc. in Engineering of Computing Systems, Politecnico di Milano.**
Graduated summa cum laude (110/110).

Awards

September 2025 **Spotlight at NeurIPS 2025.**

March 2023 **Awarded ETH AI Center Fellowship** (~ 13/650 acceptance rate).

July 2022 **Outstanding Paper Award at ICML 2022.**

July 2022 **DeepMind, ICML, and Synapse Travel Awards.**

Sept 2018 - July 2020 **Full scholarship by Politecnico di Milano for nearly perfect GPA (2 years) .**

Research Experience

March 2026 - Present **Visiting Ph.D. Student, California Institute of Technology (Caltech).**
Working on diffusion-based active exploration for new-to-nature discovery over molecular and protein design spaces. Supervised by Yisong Yue and Frances H. Arnold.

Oct 2022 - Feb 2023 **Graduate Student Researcher, University of Oxford.**
Developing theory and algorithms to introduce Group Theoretic Geometric Priors into Active Exploration Processes in Markov Chains. Supervised by prof. **Michael Bronstein** (University of Oxford, Twitter Research) and prof. **Marcello Restelli** (Politecnico di Milano).

Oct 2021 - Feb 2022 **Graduate Student Researcher, Imperial College London.**
Developed theory and algorithms to learn and exploit Causal Graphs in model-based Reinforcement Learning. Supervised by prof. **Michael Bronstein** (Imperial College London, Twitter Research) and prof. **Marcello Restelli** (Politecnico di Milano).

Jul 2019 - Feb 2021 **Undergraduate and Graduate Student Researcher, AirLab, Politecnico di Milano.**
Worked on the foundations of Unsupervised and Convex Reinforcement Learning supervised by prof. Marcello Restelli. Results of this collaboration have been awarded with the **Outstanding Paper Award at ICML 2022.**

Nov 2018 - Jul 2019 **Undergraduate Student Member, NECSTLab, Politecnico di Milano.**
Developed a software to manage biometric time series supervised by prof. Marco Domenico Santambrogio. Presented it at the Lawrence Berkeley National Laboratory and at Facebook HQ (Menlo Park, CA).

Volunteering and Mentoring

July 2022 - Present **Organizer of CS/AI Reading Group for LeadTheFuture.**
Interviewed researchers, including MIT professors, to introduce research topics such as of Computational Game Theory, Reinforcement Learning, and Domain Adaptation in ML, to a large audience of students.

Nov 2022 - Present **Mentor for LeadTheFuture.**
Mentoring italian students start their academic and research journeys.

June 2023 - Present **ETH and Caltech Student Research Supervision.**
I had the opportunity to act as thesis or project advisor for several M.S. and B.S. students, including:

- Federico Arangath (2023, ETH semester project)
- Noah Liniger (2023, ETH semester project)
- Guy Schacht (2024, ETH, M.S. thesis)
- Sven Gutjahr (2024, ETH, M.S. thesis)
- Cristian Perez Jensen (2025, ETH, M.S. thesis)
- Kimon Protopapas (2025, ETH, M.S. thesis)
- Harini Vijayshankar (2025, ETH, M.S. thesis)
- Trevor Chen (2026, Caltech, semester project)
- Ariel Dai (2026, Caltech, semester project)

Publications and Preprints

- [1] Andrey Kharitenko, Zebang Shen, **Riccardo De Santi**, Niao He, and Florian Doerfler. “Landing with the Score: Riemannian Optimization through Denoising”. In: *International Conference on Learning Representations (ICLR)*. 2026.
- [2] Cristian Perez Jensen[†], Luca Schaufelberger*, **Riccardo De Santi***, Kjell Jorner, and Andreas Krause.

- “Value Matching: Scalable and Gradient-Free Reward-Guided Flow Adaptation”. In: *International Conference on Learning Representations (ICLR)*. 2026.
- [3] **Riccardo De Santi**, Malte Franke, Ya-Ping Hsieh, and Andreas Krause. “A Unified Density Operator View of Flow Control and Merging”. In: *International Conference on Machine Learning (ICML)*, **Oral presentation** at *Workshop on Real-World Constrained and Preference-Aligned Flow and Diffusion-Based Models at ICLR 2026*. (2026).
 - [4] Sven Gutjahr^{*},[†], **Riccardo De Santi**^{*}, Luca Schaufelberger^{*}, Kjell Jorner, and Andreas Krause. “Constrained Molecular Generation via Sequential Flow Model Fine-Tuning”. In: *International Conference on Machine Learning (ICML) 2026*, **Oral presentation** at *Workshop on Frontiers in Probabilistic Inference: Sampling meets Learning at NeurIPS*. 2026.
 - [5] Zifan Wang, **Riccardo De Santi**, Xiaoyu Mo, Michael M Zavlanos, Andreas Krause, and Karl H Johansson. “Efficient Tail-Aware Generative Optimization via Flow Model Fine-Tuning”. In: *International Conference on Machine Learning (ICML)* (2026).
 - [6] **Riccardo De Santi**^{*}, Kimon Protopapas^{*}, Ya-Ping Hsieh, and Andreas Krause. “Verifier-Constrained Flow Expansion for Discovery Beyond the Data”. In: *International Conference on Learning Representations (ICLR)*. 2026.
 - [7] **Riccardo De Santi**, Marin Vlastelica, Ya-Ping Hsieh, Zebang Shen, Niao He, and Andreas Krause. “Flow Density Control: Generative Optimization Beyond Entropy-Regularized Fine-Tuning”. In: **Oral presentation** at *Workshop on Generative AI and Biology at ICML 2025*. **Spotlight** at *Conference on Neural Information Processing Systems (NeurIPS)*. 2025.
 - [8] **Riccardo De Santi**^{*}, Marin Vlastelica^{*}, Ya-Ping Hsieh, Zebang Shen, Niao He, and Andreas Krause. “Provable maximum entropy manifold exploration via diffusion models”. In: *International Conference on Machine Learning (ICML)*. 2025.
 - [9] Guy Shacht, Mojmir Mutny, **Riccardo De Santi**, Ziyad Sheebaelhamd, and Andreas Krause. “Efficient Personalization of Generative Models via Optimal Experimental Design”. In: *arXiv preprint arXiv:2512.19057*, *Workshop on Models of Human Feedback for AI Alignment at ICML 2025*. 2025.
 - [10] Mirco Mutti, **Riccardo De Santi**, Marcello Restelli, Alexander Marx^{*}, and Giorgia Ramponi^{*}. “Exploiting Causal Graph Priors with Posterior Sampling for Reinforcement Learning”. In: *International Conference on Learning Representations (ICLR) 2024*, and *Causal Rep. Learning Workshop at NeurIPS 2023*. 2024.
 - [11] **Riccardo De Santi**, Federico Arangath Joseph, Noah Liniger, Mirco Mutti, and Andreas Krause. “Geometric active exploration in Markov decision processes: the benefit of abstraction”. In: *International Conference on Machine Learning (ICML) 2024*. 2024.
 - [12] **Riccardo De Santi**^{*}, Manish Prajapat^{*}, and Andreas Krause. “Global reinforcement learning: Beyond linear and convex rewards via submodular semi-gradient methods”. In: *International Conference on Machine Learning (ICML) 2024*. 2024.
 - [13] Pietro Maldini, Mirco Mutti, **Riccardo De Santi**, and Marcello Restelli. “Recursive History Representations for Unsupervised Reinforcement Learning in Multiple-Environments”. In: *Decision Awareness in Reinforcement Learning Workshop and Pre-training: Perspectives, Pitfalls, and Paths Forward Workshop at ICML 2022*. 2023.
 - [14] Mirco Mutti, **Riccardo De Santi**, Piersilvio De Bartolomeis, and Marcello Restelli. “Convex Reinforcement Learning in Finite Trials”. In: *Journal of Machine Learning Research JMLR*. 2023.
 - [15] Mirco Mutti^{*}, **Riccardo De Santi**^{*}, and Marcello Restelli. “The Importance of Non-Markovianity in Maximum State Entropy Exploration”. In: *International Conference on Machine Learning (ICML) Outstanding Paper Award*. 2022.
 - [16] Mirco Mutti^{*}, **Riccardo De Santi**^{*}, Piersilvio De Bartolomeis, and Marcello Restelli. “Challenging Common Assumptions in Convex Reinforcement Learning”. In: *Conference on Neural Information Processing Systems (NeurIPS) 2022*, and *Complex Feedback in Online Learning Workshop at ICML 2022*. 2022.
 - [17] Mirco Mutti^{*}, **Riccardo De Santi**^{*}, Emanuele Rossi, Juan Felipe Calderon, Michael Bronstein, and Marcello Restelli. “Provably Efficient Causal Model-Based Reinforcement Learning for Systematic Generalization”. In: *In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), Workshop on Spurious Correlations, Invariance, and Stability at ICML 2022 and A Causal View on Dynamical Systems Workshop at NeurIPS 2022*. 2022.

^{*} stands for equal contribution, [†] denotes B.S. or M.S. student supervision.