

Marin Vlastelica

[Homepage](#) | [LinkedIn](#) | [Twitter](#) | [GitHub](#) | [Scholar](#)

Member of technical staff at Prometheus. My research centres on principled methods for discovery and self-improvement: agents that keep adapting and become more efficient at learning and extracting new knowledge. Background spanning reinforcement learning, generative models and LLM post-training, with experience leading research efforts in academic and industrial labs.

EDUCATION

Max Planck Institute for Intelligent Systems, Tübingen, Germany Nov 2024
PhD, Computer Science, *summa cum laude* (top 1–5%)
Thesis: *Structured, Constrained and Creative Learning*
Advisors: Georg Martius, Peter Dayan, Michael Mühlebach; examiners: Georg Martius, Peter Dayan, Mathias Niepert, Matthias Hein

Karlsruhe Institute of Technology, Karlsruhe, Germany Aug 2017
Master of Science, Artificial Intelligence (machine learning, robotics, intelligent systems)

University of Zagreb, Zagreb, Croatia Jul 2014
Bachelor of Science, Computer Science

EXPERIENCE

Member of Technical Staff, Prometheus Zurich, Switzerland · Jan 2026 – present

Research Scientist, Meta Superintelligence Labs Zurich, Switzerland · Mar 2025 – Jan 2026
Post-training for reasoning: led and coordinated efforts on data value and efficient learning in RL and SFT.
Agents: leading the self-improvement and self-play research effort for agentic models.

Guest Scientist, ETH Zurich Zurich, Switzerland · Apr 2025 – present
Student supervision on discovery with generative models and RL.

Postdoctoral Researcher, ETH AI Center Zurich, Switzerland · Oct 2024 – Apr 2025
Discovery with generative models and reinforcement learning. Hosts: Andreas Krause, Marco Hutter, Bernhard Schölkopf.

Research Scientist Intern, Meta Zurich, Switzerland · May – Oct 2024
Efficiency and fine-tuning of large-scale text-to-image and text-to-video diffusion models.

Graduate Researcher, Max Planck Institute for Intelligent Systems Tübingen, Germany · Apr 2019 – May 2024
Reinforcement learning, implicit layers and differentiable combinatorial optimization, generative models.

Research Scientist Intern, DeepMind London, United Kingdom · Jul – Dec 2022
Score-based (diffusion) generative modeling for inverse design.

Applied Research Scientist Intern, Amazon Berlin, Germany · May – Sep 2021
Reinforcement learning for task-oriented dialogue systems, amortized variational inference, normalizing flows.

Researcher, Max Planck Institute for Intelligent Systems Tübingen, Germany · Jul – Dec 2018
Hierarchical reinforcement learning.

Research Assistant, Forschungszentrum Informatik Karlsruhe, Germany · 2017
Reinforcement learning for continuous motor control via liquid state machines.

Research Intern, Volkswagen AG Data:Lab Munich, Germany · Sep – Oct 2017
Continuous control of mobile robots, path planning and unsupervised object sorting with VAEs and region proposal networks.

Research Fellow, Blue Brain Project, EPFL Geneva, Switzerland · Jan – Jul 2017
Reinforcement learning and reservoir computing for motor control with spiking neural networks.

Data Analytics Intern, IBM Research and Development Böblingen, Germany · Apr – Jul 2016
Data analytics and machine learning with Apache Spark.

Student Research Assistant, Fraunhofer IOSB Karlsruhe, Germany · Nov 2015 – Mar 2016
Knowledge modeling with Bayesian networks; built a graphical inference framework in Java.

Student Research Assistant, CVHCI, KIT Karlsruhe, Germany · 2015
Classifiers for video analysis of emotion, valence and violence (MediaEval 2015).

PUBLICATIONS

*equal contribution. Full list on [Google Scholar](#).

- Riccardo De Santi*, **Marin Vlastelica***, Ya-Ping Hsieh, Zebang Shen, Niao He, Andreas Krause, *Flow Density Control: Generative Optimization Beyond Entropy-Regularized Fine-Tuning*, NeurIPS 2025 (spotlight)
- Riccardo De Santi*, **Marin Vlastelica***, Ya-Ping Hsieh, Zebang Shen, Niao He, Andreas Krause, *Provable Maximum Entropy Manifold Exploration via Diffusion Models*, ICML 2025
- Núria Armengol Urpí, **Marin Vlastelica**, Georg Martius, Stelian Coros, *Epistemically-Guided Forward-Backward Exploration*, RLC 2025
- **Marin Vlastelica**, Jin Cheng, Georg Martius, Pavel Kolev, *Diversity Maximization under Imitation Constraints*, RLC 2024
- Jin Cheng, **Marin Vlastelica**, Pavel Kolev, Chenhao Li, Georg Martius, *Learning Diverse Skills for Local Navigation under Multi-Constraint Optimality*, ICRA 2024
- **Marin Vlastelica**, Tatiana López-Guevara, Kelsey Allen, Peter Battaglia, Arnaud Doucet, Kimberly Stachenfeld, *Diffusion Generative Inverse Design*, SPIGM workshop, ICML 2023
- Cian Eastwood*, Shashank Singh*, Andrei Liviu Nicolicioiu, **Marin Vlastelica**, Julius von Kügelgen, Bernhard Schölkopf, *Spuriousity Didn't Kill the Classifier: Using Invariant Predictions to Harness Spurious Features*, NeurIPS 2023
- Subham Sekhar Sahoo*, Anselm Paulus*, **Marin Vlastelica**, Vít Musil, Volodymyr Kuleshov, Georg Martius, *Backpropagation through Combinatorial Algorithms: Identity with Projection Works*, ICLR 2023
- **Marin Vlastelica**, Patrick Ernst, Gyuri Szarvas, *Taming Continuous Posteriors for Latent Variational Dialogue Policies*, AAAI 2023 (oral)
- Chenhao Li, **Marin Vlastelica**, Sebastian Blaes, Jonas Frey, Felix Grimmering, Georg Martius, *Learning Agile Skills via Adversarial Imitation of Rough Partial Demonstrations*, CoRL 2022 (oral, best paper nominee)
- Chenhao Li, Sebastian Blaes, Pavel Kolev, **Marin Vlastelica**, Jonas Frey, Georg Martius, *Versatile Skill Control via Self-Supervised Adversarial Imitation of Unlabeled Mixed Motions*, RSS 2022
- **Marin Vlastelica***, Sebastian Blaes*, Cristina Pinneri, Georg Martius, *Risk-Averse Zero-Order Trajectory Optimization*, CoRL 2021
- **Marin Vlastelica**, Michal Rolínek, Georg Martius, *Neuro-Algorithmic Policies Enable Fast Combinatorial Generalization*, ICML 2021 (spotlight)
- **Marin Vlastelica***, Anselm Paulus*, Vít Musil, Georg Martius, Michal Rolínek, *Differentiation of Blackbox Combinatorial Solvers*, ICLR 2020 (spotlight)
- Michal Rolínek*, Vít Musil*, Anselm Paulus, **Marin Vlastelica**, Claudio Michaelis, Georg Martius, *Optimizing Rank-Based Metrics with Blackbox Differentiation*, CVPR 2020 (best paper nominee)
- Sebastian Blaes, **Marin Vlastelica**, Jia-Jie Zhu, Georg Martius, *Control What You Can: Intrinsically Motivated Hierarchical Reinforcement Learner*, NeurIPS 2019
- Juan Camilo Vasquez Tieck, **Marin Vlastelica**, Jacques Kaiser, Arne Roennau, Marc-Oliver Gewaltig, Rüdiger Dillmann, *Learning Continuous Muscle Control for a Multi-Joint Arm by Extending Proximal Policy Optimization with a Liquid State Machine*, ICANN 2018
- **Marin Vlastelica**, Sergey Hayrapetyan, Makarand Tapaswi, Rainer Stiefelhagen, *KIT at MediaEval 2015 – Evaluating Visual Cues for Affective Impact of Movies Task*, MediaEval 2015

SERVICE

- Area Chair, European Workshop on Reinforcement Learning (EWRL) 2025
- Reviewer: NeurIPS 2020–2025, ICML 2021/2022, ICLR 2021/2022, CoRL 2021–2023, AAAI 2020/2023
- Program Chair and Organizer, Learning Meets Combinatorial Algorithms workshop, NeurIPS 2020
- Teaching Assistant, Reinforcement Learning, University of Tübingen; co-wrote a funded grant on implicit layers
- Associate Editor, Towards Data Science, 2020–present; Top Writer in Machine Learning, Medium
- Summer schools: MPC & RL (2021), CIFAR DLRL (2020), Oxford MLSS (2020), RLSS Lille (2019)

SKILLS

Research areas: reinforcement learning, LLM post-training (RL, SFT, self-play), generative and diffusion models, differentiable combinatorial optimization, exploration and discovery, optimization, statistics

Languages & frameworks: Python, C++, Java; PyTorch, JAX, TensorFlow, scikit-learn; distributed training

Open source: FoX (normalizing flows in JAX), mbri (model-based RL), torch_rl (RL in PyTorch)

Spoken languages: Croatian (native), English (C2), German (C2), Spanish (A2), Chinese (A1)

Interests: skiing, snowboarding, chess, books, basketball, guitar, piano