Lukas Schäfer

lukaschaefer.com | Google Scholar | lukas.schaefer@microsoft.com

Research Profile

I am a postdoctoral researcher at Microsoft Research, where I study how agents can learn to make decisions efficiently through online reinforcement learning and offline imitation learning. My research has appeared at NeurIPS, AAMAS, and TMLR, and I co-authored the first introductory textbook on multi-agent reinforcement learning (MIT Press, 2024), downloaded over 50,000 times.

Work Experience

Postdoctoral Researcher — Oct 2024 - Present
Microsoft Research — *Cambridge, UK*

Researched offline imitation learning algorithms for game-playing AI agents in modern video games [2, 14], and co-developed the **WHAMM world model** [15] for real-time game simulation deployed on Copilot Labs.
Advisors: Sergio Valcarcel Macua and Katja Hofmann

Research Intern — Apr 2023 - Oct 2023
Microsoft Research — *Cambridge, UK*

Empirical study on visual encoders, including pre-trained foundation models, for **game-playing agents in modern video games** [4]. **Granted funding for an extension from 3 to 6 months** following a demo to stakeholders in Xbox.
Advisors: Sam Devlin and Tabish Rashid

Research Intern — Jul 2022 - Dec 2022
Huawei Noah's Ark Lab — *London, UK*

Researched ensemble value functions to guide exploration and stabilise training in multi-agent reinforcement learning; accepted as **oral paper at AAMAS 2025** [3].
Advisor: David Mguni

Research Intern — Nov 2020 - Mar 2021
Dematic - Technology and Innovation — *Remote*

Designed multi-agent warehouse simulator and scalable learning algorithm, leading to a **fellowship-funded research collaboration**, and publication at **IROS 2024** [6].
Advisors: Aleksandar Krnjaic and Stefano V. Albrecht

Education

PhD, Data Science & Artificial Intelligence — **University of Edinburgh** — Dec 2019 - Oct 2024
Thesis: Efficient Exploration in Single-Agent and Multi-Agent Deep Reinforcement Learning
Advisors: Stefano V. Albrecht and Amos Storkey | **Award:** Passed with no corrections

MSc, Informatics — **University of Edinburgh** — Sep 2018 - Aug 2019
Thesis: Curiosity in Multi-Agent Reinforcement Learning
Advisor: Stefano V. Albrecht | **Award:** Distinction

BSc, Computer Science — **Saarland University** — Oct 2015 - Sep 2018
Thesis: Domain-Dependent Policy Learning using Neural Networks in Classical Planning
Advisor: Jörg Hoffmann | **Award:** 1.2 (within top 5%)

Publications

Citations on Google Scholar: 1,500+

Legend: Data-Driven Offline Learning Online Reinforcement Learning Multi-Agent Systems

Textbook

[1] S. V. Albrecht, F. Christianos, and **Lukas Schäfer**. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT press, 2024. **Downloaded over 50,000 times**.

Preprints

[2] **Lukas Schäfer**, A. Lemkhenter, P. Choudhury, C. Lovett, S. Nath, L. França, M. R. F. de Mendonça, A. Lamb, R. Islam, S. Sen, J. Langford, K. Hofmann, and S. V. Macua. "When does Predictive Inverse Dynamics Outperform Behavior Cloning?" In: *arxiv* (2026).

Peer-Reviewed Papers (Conferences, Journals, and Workshops)

[3] **Lukas Schäfer**, O. Slumbers, S. McAleer, Y. Du, S. V. Albrecht, and D. Mguni. "Ensemble Value Functions for Efficient Exploration in Multi-Agent Reinforcement Learning". In: *AAMAS*. 2025.

[4] **Lukas Schäfer**, L. Jones, A. Kanervisto, Y. Cao, T. Rashid, R. Georgescu, D. Bignell, S. Sen, A. T. Gavito, and S. Devlin. "Visual Encoders for Data-Efficient Imitation Learning in Modern Video Games". In: *ALA Workshop at AAMAS*. 2025.

[5] A. A. Fernandez, **Lukas Schäfer**, E. Villar-Rodriguez, S. V. Albrecht, and J. D. Ser. "Using Offline Data to Speed-up Reinforcement Learning in Procedurally Generated Environments". In: *Neurocomputing* (2024).

[6] A. Krnjaic, R. D. Steleac, J. D. Thomas, G. Papoudakis, **Lukas Schäfer**, A. W. K. To, K.-H. Lao, M. Cubuktepe, M. Haley, P. Börsting, and S. V. Albrecht. "Scalable Multi-Agent Reinforcement Learning for Warehouse Logistics with Robotic and Human Co-Workers". In: *IROS*. 2024.

[7] T. McInroe, **Lukas Schäfer**, and S. V. Albrecht. "Learning representations for control with hierarchical forward models". In: *TMLR* (2023).

[8] **Lukas Schäfer**, F. Christianos, J. P. Hanna, and S. V. Albrecht. "Decoupled Reinforcement Learning to Stabilise Intrinsically-Motivated Exploration". In: *AAMAS*. 2022.

[9] **Lukas Schäfer**. "Task Generalisation in Multi-Agent Reinforcement Learning". In: *AAMAS, DC*. 2022.

[10] R. Zhong, D. Zhang, **Lukas Schäfer**, S. V. Albrecht, and J. P. Hanna. "Robust On-Policy Data Collection for Data Efficient Policy Evaluation". In: *NeurIPS*. 2022.

[11] G. Papoudakis, F. Christianos, **Lukas Schäfer**, and S. V. Albrecht. "Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks". In: *NeurIPS, Datasets and Benchmarks*. 2021.

[12] F. Christianos, **Lukas Schäfer**, and S. V. Albrecht. "Shared Experience Actor-Critic for Multi-Agent Reinforcement Learning". In: *NeurIPS*. 2020.

[13] **Lukas Schäfer**, F. Christianos, A. Storkey, and S. V. Albrecht. "Learning Task Embeddings for Teamwork Adaptation in Multi-Agent Reinforcement Learning". In: *GenPlan Workshop at NeurIPS*. 2023.

Technical Blog Posts

[14] P. Choudhury, **Lukas Schäfer**, C. Lovett, K. Hofmann, and S. V. Macua. *Rethinking imitation learning with Predictive Inverse Dynamics Models*. Microsoft Research Blog. 2026.

[15] T. Rashid, V. Fragoso, C. Ke, Y. Cao, D. Bignell, S. Tan, **Lukas Schäfer**, S. Parisot, A. Lemkhenter, C. Lovett, P. Choudhury, R. Stevenson, S. V. Macua, A. Donnelly, D. Kennett, A. T. Gavito, L. Wen, J. Entenmann, K. Hofmann, and H. Zhang. *WHAMM! Real-time world modelling of interactive environments*. Microsoft Research Blog. 2025.

Teaching Experience

Textbook Author Mar 2022 - Dec 2024

Co-authored an introductory **textbook on multi-agent reinforcement learning** with Stefano V. Albrecht and Filippos Christianos, including the accompanying codebase and exercises for the Barcelona summer school (2024). **Downloaded over 50,000 times** since pre-release in April 2023.

Teaching Assistant Oct 2019 - Jun 2022
University of Edinburgh *Edinburgh, UK*

TA for the Reinforcement Learning course for three consecutive years; delivered lectures, designed coursework, and marked exams for 100+ students each year.

Voluntary Lecturer and Coach Sep 2017 - Oct 2017
Saarland University *Saarbrücken, Germany*

Delivered daily lectures on formal languages and logic to 250 participants in the orientation course for incoming computer science students; received the **BESTE-award** for special student commitment (Saarland University, 2017).

Teaching Assistant Oct 2016 - Mar 2017
Saarland University *Saarbrücken, Germany*

Taught weekly tutorials on functional programming, and complexity theory to undergraduates; co-designed teaching material and marked weekly tests and exams.

Supervision and Mentorship Experience

Consultation & Mentorship Sep 2025 - Present
Microsoft Research *Cambridge, UK*

Consulting engineering team to transfer recent research on vision-language models to production systems; closely mentoring a graduate-level software engineer.

Supervision of PhD Research Intern Apr 2025 - Aug 2025
Microsoft Research *Cambridge, UK*

Supervised PhD research intern Somjit Nath (McGill / Mila) on data augmentations for imitation learning in video games; insights informed research and product teams at Microsoft.

Supervision of Visiting PhD Student May 2022 - May 2023
University of Edinburgh *Edinburgh, UK*

Supervised visiting PhD student Alain Andres Fernandez (Tecnalia, Spain) on imitation learning for pre-training reinforcement learning agents during a 3-month visit; resulted in a **Neurocomputing Journal** publication [5].

Supervision of Master's Students Sep 2021 - May 2024
University of Edinburgh *Edinburgh, UK*

Supervised three Master's students during their MSc dissertation projects:

- Rahat Santosh: Integrating Agent Modelling in Multi-Agent Reinforcement Learning Algorithms
- Rujie Zhong: Data Collection for Policy Evaluation in Reinforcement Learning
  Resulted in main conference publication at **NeurIPS 2022** [10]
- Panagiotis Kyriakou: Reinforcement Learning with Function Approximation in Continuing Tasks: Discounted Return or Average Reward?

Research Community Engagement

Reviewing

- 3-time **best reviewer award** for ICML conference (2022, 2024, 2025)
- Journals: Transactions on Machine Learning Research (TMLR, 2024)
- Conferences: NeurIPS (2021, 2022, 2023), ICLR (2026), ICML (2021, 2022, 2023, 2024, 2025, 2026), AAMAS (2022, 2023, 2024), RLDM (2025), RLC (2024)
- Workshops: NeurIPS Pre-registration experiment workshop (2020)

Organisation

Lead Organiser Jul 2024 - Mar 2025
UK Multi-Agent Systems Symposium 2025 *London, UK*

Co-lead organiser of the **UK Multi-Agent Systems Symposium 2025** with 200 participants in collaboration with the Alan Turing Institute and King's College London.

Lead Organiser and Host Sep 2020 - Sep 2022
RL Reading Group, University of Edinburgh *Edinburgh, UK*

Organised and hosted RL reading group at University of Edinburgh with speakers from leading industry and academic research groups, including Google DeepMind, FAIR, Oxford University, McGill University, and National University of Singapore.

Invited Talks

Mar 2026 **University of Edinburgh**, RL & Agents Reading Group
Understanding Predictive Inverse Dynamics Models for Offline Imitation Learning

Feb 2026 **Imperial College**, I-X Seminar
Decision-Making in Modern Video GamesGames: From Game-Playing Agents to World Models

Jan 2026 **Cohere Labs Seminar**
Decision Making in Modern Video Games: From Human Play to World Models

Dec 2025 **ELLIS UnConference, Interactive Learning and Interventional Representations (ILIR) Workshop**
Exploiting State and Action Uncertainty for Imitation Learning using Inverse Dynamics Models

Oct 2025 **University of Sheffield**, ML Seminar Series
Decision Making in Modern Video Games: From Human Play to World Models

Jul 2025 **Belgium-Netherlands Workshop on Reinforcement Learning (BeNeRL)**
Decision Making in Video Games

Nov 2024 **Gazi University Turkey**, AI Research & Big Data Seminars
An Introduction to the Multi-Agent Reinforcement Learning Textbook

Mar 2024 **University of Maryland**, MARL Reading Group
An Introduction to MARL Textbook and EPyMARL Codebase

Feb 2024 **Stanford University**, Stanford Intelligent Systems Laboratory
Sample-Efficient Multi-Agent Reinforcement Learning

Jul 2022 **University of California, Berkeley**, RL Reading Group
Deep Reinforcement Learning for Multi-Agent Interaction

Skills

Natural Languages German (native), English (fluent), French (basic), Chinese (basic)
Programming Python (primary), C++, C, Bash, Java

Funding

Awarded Funding

Dec 2019 - Jun 2024	£58,731	**Principal's Career Development Scholarship** PhD scholarship from University of Edinburgh
Apr 2023 - Oct 2023	£12,860	**Microsoft Research internship extension** funding to extend research internship from 3 to 6 months; granted by Xbox after demo of internship project progress
Sep 2018 - Aug 2019	16,500 €	**DAAD graduate scholarship** postgraduate scholarship from German Academic Exchange Service
Sep 2018 - Aug 2019	£300	**Stevenson Exchange Scholarship** postgraduate scholarship to support studies in Scotland

Contributions to Funding Awards

Royal Academy of Engineering, Industrial Fellowship 2022

My internship project at Dematic initiated and laid the foundation for a multi-year research collaboration funded with **£250,000** through a **Industrial Fellowship of the Royal Academy of Engineering** between Dematic and the University of Edinburgh on scalable multi-agent reinforcement learning for warehouse logistics. The fellowship was held by my primary PhD supervisor, Prof. Stefano V. Albrecht.

Awards

Young Researcher Attendee Sep 2022
Heidelberg Laureate Forum *Heidelberg, Germany*

Selected as one of 100 international young researchers in computer science to participate in the prestigious Heidelberg Laureate Forum where I had the opportunity to network and discuss research with laureates of the most prestigious awards in mathematics and computer science.

Open-Source Software Contributions

EPyMARL – Core Contributor

Core contributor to the EPyMARL codebase, an open-source codebase for multi-agent reinforcement learning research. EPyMARL has been developed as part of a benchmarking effort [11] and has since received **over 700 stars on Github**.

MARL Textbook Codebase – Core Contributor

Core contributor to the accompanying codebase to our textbook on multi-agent reinforcement learning [1], an open-source codebase designed for ease of use and teaching in multi-agent reinforcement learning research. Within the first year of its release, the codebase has received **over 600 stars on Github**.