

Filippos Christianos

London, UK | filippos.christianos@gmail.com | fchristianos.com | linkedin.com/in/fchristianos | github.com/semitable

Summary

Research Scientist specializing in Large Language Models (LLMs) and Reinforcement Learning (RL). Proven scientific record with over 1,000 citations and publications in top venues (e.g., *NeurIPS*, *ICML*, *ICLR*). Co-Authored the multi-agent RL textbook, published by MIT Press. Strong technical background in coding and training machine learning models across diverse applications.

Experience

Research Scientist, Convergence.ai – London, UK

Jan. 2025 – present

- Post-training of Vision Language Models (VLMs) targeting web-agent control tasks, using SFT and Reinforcement Learning techniques.

Research Scientist, Huawei – London, UK

June 2023 – Jan. 2025

- Led the development of an AI-Agent (LLM based) modular framework, enabling flexible, task-specific decision-making. Achieved success rate improvements from 27% to 82% using SFT and up to 91% with RL in complex environments.
- Designed a machine learning-based control architecture for real-time task execution across mobile apps (ICLR Spotlight). Leveraged a compact Transformer with a vision-language model (VLM) to achieve up to 42% improvement over prompt-based GPT4 baselines.

Research Scientist, Intern, NVIDIA – San Francisco, California (remote)

May – Dec. 2022

- Developed BiVO, a generative model using variational autoencoders to predict occluded traffic agents' trajectories, enabling integration with planning for autonomous vehicles. Tested on the nuScenes dataset, BiVO achieved up to 8% lower planning cost compared to other occlusion predicting baselines.

Education

University of Edinburgh, PhD in Artificial Intelligence

2019 – 2023

- Used reinforcement learning algorithms (e.g., *A3C*, *PPO*, *DQN*) to develop novel multi-agent methods. Published as first author in 2x *NeurIPS*, *ICML*, and *TMLR*. Supervised by Dr Stefano Albrecht and examined by Dr Frans Oliehoek.
- Teaching Assistant for MSc-level Reinforcement Learning, designing coursework and leading sessions on Deep RL. Mentored and supervised two MSc students through their thesis projects.

University of Edinburgh, MScR in Robotics and Autonomous Systems

2018 – 2019

- Graduated with Distinction, completing coursework in Machine Learning, Reinforcement Learning, and Robotics.

Technical University of Crete, BSc & intgr. MSc (5 yrs) in Electrical and Computer Engineering

2009 – 2016

- Achieved top thesis score (10/10) and a peer-reviewed publication. Studied Computer Science, Electronics, and Telecommunications with elective courses on Artificial Intelligence, Game theory, and Multi-Agent Systems.

Selected Publications – [Google Scholar](#)

Lightweight Neural App Control

2024

First author, spotlight paper — In International Conference on Learning Representations (ICLR).

Multi-Agent Reinforcement Learning: Foundations and Modern Approaches

2024

Co-author — Textbook by The MIT Press, endorsed by Andrew Barto, Michael L. Littman, Peter Stone, and others.

Pareto Actor-Critic for Equilibrium Selection in Multi-Agent Reinforcement Learning

2023

Joint first author — In Transactions on Machine Learning Research (TMLR)

Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks

2021

Joint first author — In Advances in Neural Information Processing Systems (NeurIPS).

Scaling Multi-Agent Reinforcement Learning with Selective Parameter Sharing

2021

First author — In International Conference on Machine Learning (ICML).

Shared Experience Actor-Critic for Multi-Agent Reinforcement Learning

2020

First author — In Advances in Neural Information Processing Systems (NeurIPS)

Skills And Additional Experience

Machine Learning Frameworks: Python with PyTorch, HuggingFace, NumPy, PyTorch Lightning, and pandas.

Other Technical Skills: Expert knowledge of Linux systems and scripting. Containerization (Docker and Kubernetes). Version Control (git). Past experience with C and C++.

Technical Communication: Delivered invited lectures at University College London (2023), KTH Royal Institute of Technology (2024), RL-China (2024). Presented at seminars for institution including Berkeley Multi-Agent Seminar Group (2022), UoE (2020 – 2023), Motion2AI (2023). Skilled in presenting complex ML topics to both technical and interdisciplinary audiences.

Open-source Projects: E-PyMARL (500+ stars), MARL codebase (300+ stars), RWARE (300+ stars), LBF (150+ stars).