

# Jacob Beck

Research Scientist

Oracle  
Cambridge, Massachusetts  
+1 617 838 4009  
jakeABeck@gmail.com  
jakebeck.com  
github.com/jacooba  
Google Scholar LinkedIn

## Education

Class of 2025 **University of Oxford, United Kingdom, DPhil**, Computer Science,  
Meta-Reinforcement Learning, Multi-Agent Reinforcement Learning, In-Context Learning.  
Advisor: Prof. Shimon Whiteson

Class of 2020 **Brown University, United States, M.S.**, Computer Science, **4.0 GPA**,  
Human Feedback, Game Theory, Autonomous Vehicles, Reinforcement Learning.  
Advisor: Prof. Michael Littman, *Thesis*

Class of 2018 **Brown University, United States, B.S.**, Computer Science, **3.8 GPA**.  
Scholarships

2020 – 2024 **2020 Google DeepMind Doctoral Scholarship**

## Experience

### Full-Time

Spring 2025 **Oracle, Research Scientist**,  
– Present 

- Conducting research on agentic systems in the *Machine Learning Research Group*
- Designed pipelines for fine-tuning reasoning in coding agents with reinforcement learning
- Applied OpenAI Evolutionary Strategies for test-time training of a 20B-parameter model
- Developed a unified abstraction for agentic scaffolds, search, and evolution with LLMs.

Extended Positions

Spring 2019 **Microsoft Research, Predoc**, Advisor: Katja Hofmann, *Brown's Blog*, *MSR's Blog*,  
– Fall 2019 

- Researched long-term memory in deep RL, with first author publication at ICLR 2020
- Showed the sensitivity of modern memory approaches to stochasticity in RL
- Implemented DNC and improved over it by 9%.

Fall 2017 **Brown University Self-Driving Car Lab, Student Researcher**, Advisor: Michael Littman,  
– Fall 2020 

- Led research on learning from human demonstration with human feedback (*New Scientist*)
- Used Stackelberg game trees for human interaction, with a publication at ICSR 2019
- Created DQN to plan actions for an autonomous car in Unity.

Internships

Summer 2024 **InstaDeep, PhD Research Intern**,  

- Led research on large language models (LLMs) for protein fitness prediction
- Applied foundation models, in-context learning over related tasks, and fine-tuning
- Published the meta-learning method, *Metalic*, achieving state-of-the-art results on ProteinGym.

Fall 2018 **DeepScale, R&D Intern**,  

- Created state-of-the-art methods for lane instance segmentation using PyTorch
- Developed heuristic, cluster-based, and end-to-end approaches based on ShuffleNet.

Summer 2018 **Lyft, Software Engineer Intern,**

- Worked on behavioral planning at the Level5 autonomous vehicle lab
- Simulated human agents at a stop intersection in C++
- Coded an MDP, and solver, to find a policy for AV at a stop intersection.

Summer 2017 **Adobe, Data Science Intern,**

- Improved forecasting for the Data Science Digital Marketing research team
- Set up the models to re-train online as new data comes in (concurrently)
- Improved team's prediction accuracy by 9%, with only 121 samples per model.

Summer 2015 **Pied Piper Robotics LLC, Engineering Intern,**

- Designed, 3D printed, programmed, and wired a robot head and neck, using ROS.

## Journal Publications

2025 A Tutorial on Meta-Reinforcement Learning  
**Jacob Beck\***, Risto Vuorio\*, Evan Zheran Liu, Zheng Xiong, Luisa Zintgraf, Chelsea Finn, and Shimon Whiteson  
*Foundations and Trends in Machine Learning*

## Conference Publications

2025 Offline RLAIF: Piloting VLM Feedback for RL via SFO  
**Jacob Beck**  
*RLC Workshop on RL Beyond Rewards: Ingredients for Developing Generalist Agents*

2025 Metalic: Meta-Learning In-Context with Protein Language Models  
**Jacob Beck**, Shikha Surana, Manus McAuliffe, Oliver Bent, Thomas D Barrett, Juan Jose Garau Luis, and Paul Duckworth  
*International Conference on Learning Representations*

2024 SplAgger: Split Aggregation for In-Context Reinforcement Learning  
**Jacob Beck**, Matthew Jackson, Risto Vuorio, Zheng Xiong, and Shimon Whiteson  
*Reinforcement Learning Conference*

2024 Distilling Morphology-Conditioned Hypernetworks for Efficient Universal Morphology Control  
Zheng Xiong, Risto Vuorio, **Jacob Beck**, Matthieu Zimmer, Kun Shao, and Shimon Whiteson  
*International Conference on Machine Learning*

2023 Recurrent Hypernetworks are Surprisingly Strong in Meta-RL  
**Jacob Beck**, Risto Vuorio, Zheng Xiong, and Shimon Whiteson  
*Neural Information Processing Systems*

2023 Universal Morphology Control via Contextual Modulation  
Zheng Xiong, **Jacob Beck**, and Shimon Whiteson  
*International Conference on Machine Learning*

2023 Trust Region Bounds for Decentralized PPO Under Non-stationarity  
Mingfei Sun, Sam Devlin, **Jacob Beck**, Katja Hofmann, and Shimon Whiteson  
*Autonomous Agents and Multiagent Systems (Best Paper Award)*

2022 Hypernetworks in Meta-Reinforcement Learning  
**Jacob Beck**, Matthew Jackson, Risto Vuorio, and Shimon Whiteson  
*Conference on Robot Learning*

2021 On the Practical Consistency of Meta-Reinforcement Learning Algorithms  
Zheng Xiong, Luisa Zintgraf, **Jacob Beck**, Risto Vuorio, and Shimon Whiteson  
*NeurIPS Meta-Learning Workshop*

---

\*Contributed equally

2021 No DICE: An Investigation of the Bias-Variance Trade off in Meta-Gradients  
Risto Vuorio, **Jacob Beck**, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson  
*NeurIPS Deep RL Workshop*

2020 Stackelberg Punishment and Bully-Proofing Autonomous Vehicles  
Matt Cooper, Jun Ki Lee, **Jacob Beck**, Joshua D. Fishman, Michael Gillett, Zoë Papakipos, Aaron Zhang, Jerome Ramos, Aansh Shah, and Michael L. Littman  
*International Conference on Social Robotics*

2020 AMRL: Aggregated Memory For Reinforcement Learning  
**Jacob Beck**, Kamil Ciosek, Sam Devlin, Sebastian Tschiatschek, Cheng Zhang, and Katja Hofmann  
*International Conference on Learning Representations*

## Academic Services

Speaker

2024 **ICML Panel, Auto-RL, Panelist**,  
Invited to speak on a *workshop panel* about meta-learning, auto-ML, and LLMs at ICML.

2024 **Meta-RL Tutorial, Presenter**,  
Presented a *tutorial* on meta-reinforcement learning at AAAI 2024.

2023 **TalkRL Podcast, Interview**,  
Interviewed on the *TalkRL Podcast* to explain meta-reinforcement learning.

2023 **Meta-RL Tutorial, Presenter**,  
Presented a *tutorial* on meta-reinforcement learning at AutoML 2023.

2021 **ICLR Panel, Philosophy and AGI, Organizer**,  
Organized two *panels* on the the technical and philosophical problems of AGI.

Reviewer

2025 **Transactions on Machine Learning Research**, *Reviewer*.

2025 **Reinforcement Learning Conference**, *Reviewer*.

2025 **International Conference on Learning Representations**, *Reviewer*.

2023 **Neural Information Processing Systems**, *Top Reviewer*.

2022 **NeurIPS Deep RL Workshop**, *Reviewer*.

2022 **International Conference on Machine Learning**, *Reviewer*.

2020 **Nature Communications**, *Reviewer*.

## Teaching and Supervision

Spring, 2022 **Supervisor**, *University of Oxford*,  
Matthew Jackson, Hypernetworks in meta-reinforcement learning.

Fall, 2021 **Co-Supervisor**, *University of Oxford*,  
Zheng Xiong, On the practical consistency of meta-reinforcement learning algorithms.

Fall, 2017 **Deep Learning Teaching Assistant**, *Brown University*,  

- Designed, taught, graded material for the graduate deep learning course, CSCI 2470
- Gave a lecture on implementing sequence-to-sequence translation with attention
- Designed a lab on recurrent neural networks.