Alexander D. Goldie

alexdgoldie@gmail.com

GitHub · Google Scholar · LinkedIn

About Me

I am a final-year PhD student at the *University of Oxford*. My goal is to accelerate research in AI through autonomous algorithm discovery using large language models, meta-learning and evolutionary methods. I am also passionate about reinforcement learning and creating more open-ended AI systems.

Education

University of Oxford – DPhil in Engineering Science **2022–2026**

Member of the AIMS CDT, funded by EPSRC.

Supervised by Jakob N. Foerster (FLAIR) and Shimon Whiteson (WhiRL).

University of Oxford – MEng in Engineering Science **2018–2022**

Overall grade: **First-Class Honours (78%)**.

Master's project supervised by Janet Pierrehumbert.

Selected modules: Machine Learning, Convex Optimization, Advanced Probability.

Publications

[1] **[Outstanding Paper] Alexander D. Goldie**, Zilin Wang, Jaron Cohen, Jakob N. Foerster, and Shimon Whiteson. How Should We Meta-Learn Reinforcement Learning Algorithms? In *Reinforcement Learning Conference*, 2025. Awarded 'Outstanding Paper For Scientific Understanding In Reinforcement Learning'.

[2] **[Spotlight] Alexander D. Goldie**, Chris Lu, Matthew T. Jackson, Shimon Whiteson, and Jakob N. Foerster. Can Learned Optimization Make Reinforcement Learning Less Difficult? In *Advances in Neural Information Processing Systems*, 2024. Also awarded a **spotlight** in the AutoRL Workshop At ICML24.

[3] Nathan Monette, Alistair Letcher, Michael Beukman, Matthew Thomas Jackson, Alexander Rutherford, **Alexander D. Goldie**, and Jakob Nicolaus Foerster. An Optimisation Framework for Unsupervised Environment Design. In *Reinforcement Learning Conference*, 2025.

[4] Benjamin Ellis*, Matthew T. Jackson*, Andrei Lupu, **Alexander D. Goldie**, Mattie Fellows, Shimon Whiteson, and Jakob N. Foerster. Adam on Local Time: Addressing Nonstationarity in RL with Relative Adam Timesteps. In *Advances in Neural Information Processing Systems*, 2024.

[5] Uljad Berdica, Kelvin Li, Michael Beukman, **Alexander D. Goldie**, Mattie Fellows, Perla Maiolino, and Jakob N. Foerster. Robust Offline Learning via Adversarial World Models. In *Open-World Agents Workshop & Adversarial ML Workshop at NeurIPS24*, 2024.

Experience

Wayve – Research Scientist Intern, Offline RL Team **Jun–Nov 2025**

Implemented a distributional reinforcement learning system, leading to a **statistically significant** driving improvement in every geography Wayve operates in.

Integrated prioritised experience replay into the distributed data pipeline to explore more data-efficient learning, and tested novel data prioritisation metrics.

Improved optimisation stability for training large (multi-billion parameter) driving models with offline reinforcement learning.

Pierrehumbert Language Modeling Group – Master's Project **2021–2022**

Developed a directed similarity metric, based on text embeddings, to investigate linguistic similarity between scientific fields.

Researched directed link prediction on the resultant linguistic graph.

Thesis awarded high first-class grade (**78%**).

Selected Awards

Outstanding Paper for Scientific Understanding in RL **2025**
Reinforcement Learning Conference

Spotlight Paper **2024**
NeurIPS Main Conference

Spotlight & Contributed Talk **2024**
ICML AutoRL Workshop

Outstanding Reviewer **2024**
ICML AutoRL Workshop

Hilary Haworth Prize (Best 2nd/3rd Year Scientist) **2021**
St. Hugh's College

Academic Service

Supervision: Theo Wolf (CDT, Ongoing), Albert Zuo (MSc, Distinction), Jaron Cohen (MSc, Distinction), Nathan Monette (Visiting, RLC Paper)

Teaching: Reinforcement Learning (PhD Course), Machine Learning (Master's Course)

Reviewer: ICML (2025 Conference, 2024 AutoRL Workshop (**award**), 2023 TiFA Workshop); ICLR (2024 Conference); NeurIPS (2025 & 2024 Conference, 2023 SoLaR Workshop)

Panelist: ICML AutoRL Workshop 2024

Podcast: Featured in two episodes of TalkRL.

Outreach: Oxford Maths Festival (**People's Choice Award**)