

Anya Sims

📍 University of Oxford ✉ anya.sims@stats.ox.ac.uk 🔗 anyasims.github.io 🌐 anyasims

Education

University of Oxford Sept 2022 – 2026

PhD in Machine Learning and Modern Statistics

Supervisors: Yee Whye Teh (OxCSML Group, and Research Director at Google DeepMind)
 Jakob N. Foerster (FLAIR Lab Lead, and Meta AI)

University of Cambridge Sept 2018 – 2022

MEng in Information and Computer Engineering

BEng First Class with Distinction (ranked 5/288)

Department prize for best 3rd year project

Top-ranked final exam performance in Information Engineering Instrumentation and Control

MEng First Class with Distinction (ranked 8/261)

Received Distinction for dissertation – Supervised by Phil Woodland

Several Engineering Department and Jesus College Cambridge awards/scholarships each year

Research Experience

Current main project

Adversarial, Open-Ended RL Training for LLMs

Adversarial RL for LLMs on dynamically generated, open-ended, yet verifiable code tasks.

As part of FLAIR, Oxford.

First Author Papers

Fill-In-The-Gaps Reinforcement Learning for Mid-Training on Unstructured Text Mar 2026

Eltayeb Ahmed*, **Anya Sims***, Thomas Foster, Tim Rocktäschel, Jakob N. Foerster

Under review for RLC 2026.

StochasTok: Improving Fine-Grained Subword Understanding in LLMs [ArXiv] June 2025

Anya Sims, Thom Foster, Klara Kaleb, Tuan-Duy H. Nguyen, Joseph Lee, Jakob N. Foerster, Yee Whye Teh, Cong Lu

Accepted at ICLR 2026. Oral at BuildingTrust ICLR 2025 Workshop.

Reinforcing General Reasoning without Verifiers [ArXiv] May 2025

Xiangxin Zhou*, Zichen Liu*, **Anya Sims***, Haonan Wang, Tianyu Pang, Chongxuan Li, Liang Wang, Min Lin, Chao Du

Accepted at ICLR 2026. Spotlight at MTI-LLM NeurIPS 2025 Workshop.

GEM: A Gym for Agentic LLMs [ArXiv] May 2025

Zichen Liu*, **Anya Sims***, Keyu Duan*, Changyu Chen*, Simon Yu, Xiangxin Zhou, Haotian Xu, Shaopan Xiong, Bo Liu, Chenmien Tan, Chuen Yang Beh, Weixun Wang, Hao Zhu, Weiyang Shi, Diyi Yang, Michael Shieh, Yee Whye Teh, Wee Sun Lee, Min Lin

Accepted at ICLR 2026.

The Edge-of-Reach Problem in Offline Model-Based Reinforcement Learning [ArXiv] Feb 2024

Anya Sims, Cong Lu, Jakob N. Foerster, Yee Whye Teh

In proceedings at NeurIPS 2024.

Core Contributor Papers

Evolution Strategies at the Hyperscale [ArXiv] Nov 2025

Bidipta Sarkar*, Mattie Fellows*, Juan Agustin Duque*, Alistair Letcher[†], Antonio León Villares[†], **Anya Sims[†]**, Dylan Cope[†], Jarek Liesen[†], Lukas Seier[†], Theo Wolf[†], Uljad Berdica[†], Alexander David Goldie, Aaron Courville, Karin Sevegnani, Shimon Whiteson, Jakob Nicolaus Foerster

Under review at ICML 2026.

- LILO: Learning to Reason at the Frontier of Learnability** [ArXiv] Feb 2025
 Thomas Foster, **Anya Sims**, Johannes Forkel, Mattie Fellows, Jakob Foerster
In proceedings at NeurIPS 2025.
- Other Papers and Research Projects*
- MARCOS: Deep Thinking by Markov Chain of Continuous Thoughts** [ArXiv] Sep 2025
 Jiayu Liu, Zhenya Huang, **Anya Sims**, Enhong Chen, Yee Whye Teh, Ning Miao
Under review at ICML 2026.
- Fine-tuning LLMs using gradient-designed direct preference surrogate objectives**
 CDT Project supervised by Yee Whye Teh and Jakob N. Foerster, with Cong Lu, and Silvia Sapora.
- Investigating plasticity loss and the effects of non-stationary training in deep neural networks**
 CDT Project supervised by Yee Whye Teh and Chris Gamble (Google DeepMind).
- Sphering transformations for stable equivariant normalization in graph neural networks**
 CDT Project supervised by Yee Whye Teh and with Sheheryar Zaidi.
- Identifiable variational autoencoders for disentangled causal representation learning** Summer 2022
Research Placement Project – University of Cambridge
 Supervised by José Miguel Hernández-Lobato and with Wenlin Chen.
- Efficient Decoding in End-to-End Speech Recognition in Transducer Language Models** 2021-22
MEng Dissertation – University of Cambridge
 Supervised by Phil Woodland, with Tony Zheng.

Industry Experience

- Associate Member, Sea AI Lab** Apr 2025 – Jan 2026
 ◦ Working across several projects, including a comprehensive codebase of multi-task, multi-turn language environments for developing and scaling multi-turn RL algorithms for training LLMs with tool capabilities.
- Software Engineering Intern, Graphcore** Summer 2021
 ◦ Worked within the AI Applications team to implement a codebase demonstrating several image recognition benchmarks on Graphcore’s IPU in TensorFlow2. This involved coordinating loading of different stages of the model onto different IPU processing cores to maximize training and inference efficiency.
- Machine Learning Mentorship, Featurespace** Summer 2020
 ◦ Self-led investigation of different deep learning methods on the Kaggle Abstraction and Reasoning Challenge.
- Software Engineering Intern, Softwire** Dec 2019
 ◦ Website development – Worked primarily on asynchronous API calls component (JavaScript and React).
- Machine Learning Intern, Infosys** Summer 2019
 ◦ Reproduced results from several recent papers on aspect-based sentiment analysis (TensorFlow and Keras).

Other

- Teaching:** **Advanced Topics in Statistical Machine Learning** 2023 – 2024
TA and Tutor, 4th Year Maths and MSc course
 SVMs, kernel methods, Bayesian machine learning, Gaussian processes, Bayesian optimization, deep learning, latent variable models.
- Languages/Frameworks:** Python, C++, Bash, PyTorch, JAX, vLLM, verl
- Roles:** Cambridge University Windsurfing Secretary (2020-2022)
 Volunteer secondary school maths tutor (2018-2019)