

Cong Lu

conglu.co.uk | conglu97@outlook.com | Google Scholar | GitHub

Research Appointments

Google DeepMind

Research Scientist, Open-Endedness Team and Gemini Team

London, UK

Feb 2025–Present

- Lead of a 15-person team focused on self-improvement and evaluation for Gemini models and agents.
- Contributed to public-facing releases: *Genie 3*, a world model that generates an unprecedented diversity of interactive environments in real time and was selected for *TIME Best Inventions of 2025 (Immersive Technology)*, and *SIMA 2*, a Gemini-powered agent that can play, reason, and learn in virtual 3D worlds.

Sakana AI

Visiting Research Scientist

Remote

Sep 2024–Feb 2025

- Released *The AI Scientist*, the first agent to automate the entire scientific process: from forming hypotheses and conducting experiments to visualizing results, writing a paper, and reviewing it.

University of British Columbia and Vector Institute

Postdoctoral Research and Teaching Fellow

Vancouver, Canada

Jan 2024–Feb 2025

- Worked in the UBC AI Research Lab with Prof. Jeff Clune on open-endedness and autonomous scientific discovery, including *The AI Scientist*, *Intelligent Go-Explore*, *Automated Capability Discovery*, and *Darwin Godel Machine*.

Waymo Research

Research Intern

Oxford, UK

Aug 2022–Dec 2023

- Worked on controllable driving trajectory generation with diffusion models for rare tail-scenario generation in autonomous driving simulation.

Microsoft Research

Research Intern

Cambridge, UK

Apr 2022–Jul 2022

- Collaborated with Microsoft Xbox Game Studios on scalable AI methods for testing video games, leading to *Go-Explore Complex 3-D Game Environments for Automated Reachability Testing (IEEE Transactions on Games, 2024)*.

Education

University of Oxford

Doctor of Philosophy in Autonomous Intelligent Machines and Systems

Oxford, UK

2019–2023

Thesis

Towards Efficient and Robust Reinforcement Learning via Synthetic Environments and Offline Data

Advisors

Prof. Michael A. Osborne and Prof. Yee Whye Teh

- Member of the Oxford Computational Statistics and Machine Learning (OxCSML) group and the Machine Learning Research Group (MLRG).

University of Oxford

Master of Mathematics and Computer Science

Oxford, UK

2015–2019

Thesis

Reinforcement Learning with Corrupted or Noisy Reward

Advisor

Prof. Varun Kanade

- Graduated with First Class Honours (highest classification) in all four years.

- Awarded Les Wood Scholarship for academic performance (2017–2019).
- Awarded Gibbs Prize in Computer Science for highest exam performance (**ranked 1st in year**).

Awards and Prizes

2025	<i>TIME Best Inventions of 2025 (Immersive Technology)</i> , Google DeepMind <i>Genie 3</i>
2024	Oral Presentation, Language Gamification Workshop @ NeurIPS
2024	Outstanding Paper Award (top 3/103), Open World Agents Workshop @ NeurIPS
2024	Outstanding Paper Award, GenAI4DM Workshop @ ICLR
2022	Outstanding Reviewer Award, NeurIPS
2022	Best Paper Award, Workshop on Learning from Diverse, Offline Data @ RSS
2018	First Place, C1 Data Open Championship (\$20K cash prize)
2017–2019	Les Wood Scholarship, Balliol College, University of Oxford
2016–2018	Gibbs Prize for Computer Science in Parts A and B, University of Oxford (ranked 1st)
2016	Les Wood Exhibitioner, Balliol College, University of Oxford

Grants Won

2025	Killam Accelerator Research Fellowship (\$110K CAD, co-writer)
2024	Center for AI Safety Compute Cluster Grant (approximately \$300,000 USD, co-writer)

Selected Preprints and Industry Releases

* indicates shared first authorship; † indicates shared senior authorship.

-
- [1] Ball, Philip J., ... , **Cong Lu**, ... *Genie 3: A New Frontier for World Models*. DeepMind Blog, 2025.
 - [2] Bolton, Andrew, ... , **Cong Lu**, ... *SIMA 2: A Generalist Embodied Agent for Virtual Worlds*. DeepMind Blog, 2025.
 - [3] Chris Lu*, **Cong Lu***, Robert Tjarko Lange*, Jakob Foerster†, Jeff Clune†, and David Ha†. *The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery*. Preprint (arXiv:2408.06292), 2024.

Peer-Reviewed Journal and Conference Publications

- [1] Chris Lu*, **Cong Lu***, Robert Tjarko Lange*, Yutaro Yamada*, Shengran Hu, Jakob Nicolaus Foerster, David Ha, and Jeff Clune. *Towards End-to-End Automation of AI Research*. *Nature*, 651:914–919, March 2026.
- [2] Anya Sims, Thomas Foster, T. Duy Nguyen-Hien, Klara Kaleb, Joseph Lee, Jakob Nicolaus Foerster, Yee Whye Teh, and **Cong Lu**. *StochasTok: Improving Fine-Grained Subword Understanding in LLMs*. In *ICLR*, 2026.
- [3] Jenny Zhang*, Shengran Hu*, **Cong Lu**, Robert Tjarko Lange†, and Jeff Clune†. *Darwin Godel Machine: Open-Ended Evolution of Self-Improving Agents*. In *ICLR*, 2026.
- [4] Aaron Dharna, **Cong Lu**, and Jeff Clune. *Foundation Model Self-Play: Open-Ended Strategy Innovation via Foundation Models*. In *Reinforcement Learning Conference*, 2025.
- [5] Clémence Grislain, Risto Vuorio, **Cong Lu**, and Shimon Whiteson. *IGDrivSim: A Benchmark for the Imitation Gap in Autonomous Driving*. In *IROS*, 2025.
- [6] Shengran Hu, **Cong Lu**, and Jeff Clune. *Automated Design of Agentic Systems*. In *ICLR*, 2025.
- [7] Nishant Jha*, Joshua Kravitz*, Jacob West-Roberts*, **Cong Lu**, Antonio Pedro Camargo, Simon Roux, Andre Cornman, and Yunha Hwang. *Gaia: An AI-Enabled Genomic Context-Aware Platform for Protein Sequence Annotation*. *Science Advances*, 11(25):eadv5109, 2025.

- [8] **Cong Lu**, Shengran Hu, and Jeff Clune. Intelligent Go-Explore: Standing on the Shoulders of Giant Foundation Models. In *ICLR*, 2025.
- [9] Gunshi Gupta*, Karmesh Yadav*, Yarin Gal, Dhruv Batra, Zsolt Kira, **Cong Lu**, and Tim G. J. Rudner. Pre-Trained Text-to-Image Diffusion Models Are Versatile Representation Learners for Control. In *NeurIPS (Spotlight)*, 2024.
- [10] Matthew T. Jackson, Michael T. Matthews, **Cong Lu**, Benjamin Ellis, Shimon Whiteson[†], and Jakob Nicolaus Foerster[†]. Policy-Guided Diffusion. In *Reinforcement Learning Conference (Oral)*, 2024.
- [11] Andrew Melnik, Michal Ljubljanac, **Cong Lu**, Qi Yan, Weiming Ren, and Helge Ritter. Video Diffusion Models: A Survey. *TMLR*, 2024.
- [12] Anya Sims, **Cong Lu**, Jakob Nicolaus Foerster, and Yee Whye Teh. The Edge-of-Reach Problem in Offline Model-Based Reinforcement Learning. In *NeurIPS*, 2024.
- [13] **Cong Lu**, Raluca Georgescu[†], and Johan Verwey[†]. Go-Explore Complex 3-D Game Environments for Automated Reachability Testing. *IEEE Transactions on Games*, 16(1):235–240, 2024.
- [14] **Cong Lu***, Philip J. Ball*, Tim G. J. Rudner, Jack Parker-Holder, Michael A. Osborne, and Yee Whye Teh. Challenges and Opportunities in Offline Reinforcement Learning from Visual Observations. *TMLR*, 2023.
- [15] **Cong Lu***, Philip J. Ball*, Yee Whye Teh, and Jack Parker-Holder. Synthetic Experience Replay. In *NeurIPS*, 2023.
- [16] **Cong Lu***, Philip J. Ball*, Jack Parker-Holder, Michael A. Osborne, and Stephen J. Roberts. Revisiting Design Choices in Offline Model-Based Reinforcement Learning. In *ICLR (Spotlight)*, 2022.
- [17] Xingchen Wan, **Cong Lu**, Jack Parker-Holder, Philip J. Ball, Vu Nguyen, Binxin Ru, and Michael A. Osborne. Bayesian Generational Population-Based Training. In *AutoML*, 2022.
- [18] Philip J. Ball*, **Cong Lu***, Jack Parker-Holder, and Stephen Roberts. Augmented World Models Facilitate Zero-Shot Dynamics Generalization From a Single Offline Environment. In *ICML*, 2021.
- [19] Tim G. J. Rudner*, **Cong Lu***, Michael A. Osborne, Yarin Gal, and Yee Whye Teh. On Pathologies in KL-Regularized Reinforcement Learning from Expert Demonstrations. In *NeurIPS*, 2021.
- [20] Xingchen Wan, Vu Nguyen, Huong Ha, Binxin Ru, **Cong Lu**, and Michael A. Osborne. Think Global and Act Local: Bayesian Optimisation over High-Dimensional Categorical and Mixed Search Spaces. In *ICML*, 2021.
- [21] Luisa Zintgraf, Leo Feng, **Cong Lu**, Maximilian Igl, Kristian Hartikainen, Katja Hofmann, and Shimon Whiteson. Exploration in Approximate Hyper-State Space for Meta Reinforcement Learning. In *ICML*, 2021.
- [22] Luisa Zintgraf, Sebastian Schulze, **Cong Lu**, Leo Feng, Maximilian Igl, Kyriacos Shiarlis, Yarin Gal, Katja Hofmann, and Shimon Whiteson. VariBAD: Variational Bayes-Adaptive Deep RL via Meta-Learning. *JMLR*, 2021.

Invited Talks

Jun 2025	Meta AI, <i>Towards Fully Autonomous Open-Ended Scientific Discovery</i>
Apr 2025	Workshop on Advances of Generative AI @ ICLR 2025 (Keynote), <i>Towards Fully Autonomous Open-Ended Scientific Discovery</i>
Mar 2025	AI4Research Workshop @ AAI 2025 (Keynote), <i>Knowledge-Grounded Autonomous Scientific Research</i>
Feb 2025	Bank of England, <i>The AI Scientist</i>
Dec 2024	Google DeepMind, <i>Towards Fully Autonomous Open-Ended Scientific Discovery</i>
Nov 2024	University of Tokyo, <i>The AI Scientist</i>
Oct 2024	Cohere For AI, <i>The AI Scientist</i>

Sep 2024	Huawei Research, <i>Automating Scientific Discovery in the Age of Foundation Models</i>
Sep 2024	Microsoft Research, <i>The AI Scientist</i>
Sep 2024	Shanghai AI Lab, <i>The AI Scientist</i>
Sep 2023	Towards Autonomous Robotic Systems Conference UK, <i>Synthetic Experience Replay</i>

Teaching Experience

Fall 2024	Course Instructor, CS340/540 Machine Learning and Data Mining, University of British Columbia (UBC) (36 lectures, cross-listed at undergraduate and graduate levels). Enrolled 258 students, managed 8 TAs. Overall rating: 4.3/5.0 (91% approval rating).
2022	Teaching Assistant, Advanced Simulation (Statistics, Oxford)
2021	Teaching Assistant, Imperative Programming (Computer Science, Oxford)
2021	Teaching Assistant, Probability, Measure and Martingales (Mathematics, Oxford)

Academic Service

Reviewing

- Journals: Nature Machine Intelligence
- Conferences: AISTATS (2021), ICML (2022–2024), NeurIPS (2022–2024; top 8% of reviewers in 2022), ICLR (2024–2025)
- Workshops: ICLR Tiny Papers (2023), Reincarnating RL Workshop @ ICLR (2023), NeurIPS MINT Workshop (2024), ICLR RSI Workshop (2026)

Program Committee

- Foundation Models for Decision Making Workshop @ NeurIPS (2022–2023)
- RL for Real Life Workshop @ NeurIPS (2022)
- Agent Learning in Open-Endedness Workshop @ ICLR (2022) and NeurIPS (2023)

Supervision

- Hannah Erlebach (Master's), University College London (UCL) Master's Thesis. Achieved distinction and is now a PhD candidate at the University of Oxford.
- Clémence Grislain (Master's), University of Oxford Summer Research Internship (2023). This work was published at IROS, and she is now a PhD candidate at Sorbonne University.
- Aneta Swianiewicz (Master's), University of Oxford UNIQ+ research internship with DeepMind and final-year project (2022–2023). Achieved first-class honours with 85% in the dissertation.
- Theodore-Mihai Iliant (Undergraduate), Oxford Computational Statistics and Machine Learning (OxCSML) Predoctoral Research Opportunities Program (8 weeks, Summer 2021).

Selected Media Coverage and Interviews

Mar 2026	Nature Editorial, <i>AI scientists are changing research — institutions, funders and publishers must respond</i>
Mar 2026	Nature News, <i>How to build an AI scientist: first peer-reviewed paper spills the secrets</i>
Mar 2026	Scientific American, <i>AI wrote a scientific paper that passed peer review</i>
Jan 2026	Reuters, <i>Videogame stocks slide on Google's AI model that turns prompts into playable worlds</i>
Nov 2025	The Verge, <i>Watch Google DeepMind's new AI agent learn to play video games</i>
Nov 2025	TechCrunch, <i>Google's SIMA 2 agent uses Gemini to reason and act in virtual worlds</i>

Oct 2025	TIME, <i>Google DeepMind Genie 3 (Best Inventions of 2025, Immersive Technology)</i>
Oct 2025	State of AI Report, featured coverage of <i>The AI Scientist</i>
Aug 2025	The Guardian, <i>Google says its new 'world model' could train AI robots in virtual warehouses</i>
Aug 2025	TechCrunch, <i>DeepMind thinks its new Genie 3 world model presents a stepping stone toward AGI</i>
Aug 2025	The Verge, <i>Google's new AI model creates video game worlds in real time</i>
Jun 2025	Fortune, <i>AI that can modify and improve its own code is here. Does this mean OpenAI's Sam Altman is right about the singularity?</i>
Mar 2025	ML Street Talk, interview on <i>The AI Scientist</i>
Feb 2025	CBC's <i>Quirks & Quarks</i> , interview on how AI is transforming science
Oct 2024	State of AI Report, featured coverage of <i>The AI Scientist</i>
Sep 2024	IEEE Spectrum, <i>Will the "AI Scientist" Bring Anything to Science?</i>
Aug 2024	VentureBeat, <i>Sakana AI's 'AI Scientist' conducts research autonomously, challenging scientific norms</i>
Aug 2024	WIRED, <i>An 'AI Scientist' Is Inventing and Running Its Own Experiments</i>
Aug 2024	Nature News, <i>Researchers built an 'AI Scientist' — what can it do?</i>

Skills

Programming	C, C++, Python, Rust, Java, Scala, Haskell
ML Frameworks	PyTorch, TensorFlow, JAX
Languages	English (native), Chinese (native)