

Yingchen Xu

✉ yingchen.xu.21@ucl.ac.uk | 🏠 ycxuyingchen.github.io | 🐦 @YingchenX

Education

University College London

PHD COMPUTER SCIENCE

- Advisors: Dr. Tim Rocktäschel, Dr. Edward Grefenstette
- Research Interests: Deep Reinforcement Learning, Embodied AI, LLM Reasoning

London, UK

Jan 2022 - present

Rice University

BS COMPUTER SCIENCE, BA MATHEMATICS

- Undergraduate Research Advisor: Dr. Anshumali Shrivastava

Houston, TX, USA

Aug 2014 - May 2018

Professional Experience

Sakana AI

RESEARCH INTERN

- Building biologically inspired data efficient visual world models.

Tokyo, Japan

Sep 2025 - Present

AI at Meta

PHD STUDENT RESEARCHER

- Conduct fundamental research in deep reinforcement learning and generative AI models.
- First-authored publications at NeurIPS and ICLR.

London, UK

Jan 2022 - Sep 2025

Airbnb

MACHINE LEARNING ENGINEER

- Built ML models for dynamic pricing and personalized content recommendations.
- Developed in-house ML infrastructures for feature engineering, model training and deployment.

San Francisco, CA, USA

Aug 2018 - Sep 2021

Publications

Zhengyao Jiang*, **Yingchen Xu***, Nolan Wagener, Yicheng Luo, Michael Janner, Edward Grefenstette, Tim Rocktäschel, Yuandong Tian. H-GAP: Humanoid Control with a Generalist Planner. *ICLR Spotlight*, 2024.

Rohan Chitnis*, **Yingchen Xu***, Bobak Hashemi, Lucas Lehnert, Urun Dogan, Zheqing Zhu, Olivier Delalleau. IQL-TD-MPC: Implicit Q-Learning for Hierarchical Model Predictive Control. *ICRA*, 2024.

Yingchen Xu*, Jack Parker-Holder*, Aldo Pacchiano*, Philip J. Ball*, Oleh Rybkin, Stephen J. Roberts, Tim Rocktäschel, Edward Grefenstette. Learning General World Models in a Handful of Reward-Free Deployments. *NeurIPS*, 2022.

Andrea Tirinzoni1*, Ahmed Touati*, Jesse Farebrother, Mateusz Guzek, Anssi Kanervisto, **Yingchen Xu**, Alessandro Lazaric, Matteo Pirotta. Zero-Shot Whole-Body Humanoid Control via Behavioral Foundation Models. *ICLR*, 2025.

Harshit Sikchi, Andrea Tirinzoni, Ahmed Touati, **Yingchen Xu**, Anssi Kanervisto, Scott Niekum, Amy Zhang, Alessandro Lazaric, Matteo Pirotta. Fast Adaptation with Behavioral Foundation Models. *RLC*, 2025.

Andy Su, Hanlin Zhu*, **Yingchen Xu***, Jiantao Jiao, Qinqing Zheng**, Yuandong Tian**. Token Assorted: Mixing Latent and Text Tokens for Improved Language Model Reasoning. *ICML*, 2025.

Beidi Chen, **Yingchen Xu**, Anshumali Shrivastava. LGD: Fast and Accurate Stochastic Gradient Estimation. *NeurIPS*, 2019.

Ari Beller, **Yingchen Xu**, Scott Linderman, Tobias Gerstenberg. Looking into the past: Eye-tracking mental simulation in physical inference. *Cognitive Science*, 2022.

Additional Experience

- 2025 **Workshop Organizer**, Reinforcement Learning Beyond Rewards Workshop at RLC
- 2023 **Workshop Organizer**, Agent Learning in Open-Endedness (ALOE) Workshop at NeurIPS
- 2023-2025 **Mentor**, MSc Research Project in Offline Reinforcement Learning, University College London
- 2022-2025 **Reviewer**, NeurIPS, ICLR and ICML