

Kurtland Chua

chua.kurtland@gmail.com • [kchua.github.io](https://github.com/kchua) • (858) 226-4053

Education

Princeton University

Ph.D. Computer Science (GPA: 4.00/4.00)

2019 - Present

Supported by a National Science Foundation Graduate Research Fellowship (NSF-GRFP)

University of California, Berkeley

B.S. Electrical Engineering and Computer Sciences (GPA: 4.00/4.00)

2015 - 2019

Graduated with Highest Honors, Phi Beta Kappa, EECS Major Citation

Experience

Princeton University, Graduate Student Researcher

09/2019 - Present

Theoretically analyzed meta-learning/few-shot learning in supervised and reinforcement learning (RL) settings; explored a multimodal approach (e.g. using natural language guidance) to meta-learning.

Facebook Artificial Intelligence Research, Research Intern

06/2019 - 08/2019

Developed a method for inferring hierarchies from data to enable model-based RL with long planning horizons. Collaborated on a separate project which demonstrated the benefit of automatic and dynamic hyperparameter tuning for model-based RL over human experts.

Berkeley AI Research Laboratory, Undergraduate Researcher

09/2016 - 05/2019

Explored using uncertainty estimation in deep model-based RL for data efficiency, creating a state-of-the-art algorithm requiring 8x fewer samples than prior methods. Solely developed and open-sourced the [project code-base](#), which has been widely used by the research community since.

Selected Publications

- Kurtland Chua, Qi Lei, Jason D. Lee. (2023) "Provable Hierarchy-Based Meta-Reinforcement Learning." *International Conference on Artificial Intelligence and Statistics (AISTATS) 2023*.
- Kurtland Chua, Qi Lei, Jason D. Lee. (2021) "How Fine-Tuning Allows for Effective Meta-Learning." *Neural Information Processing Systems (NeurIPS) 2021*.
- Baohe Zhang, Raghu Rajan, Luis Pineda, Nathan Lambert, André Biedenkapp, Kurtland Chua, Frank Hutter, Roberto Calandra. (2021) "On the Importance of Hyperparameter Optimization for Model-Based Reinforcement Learning." *International Conference on Artificial Intelligence and Statistics (AISTATS) 2021*.
- Kurtland Chua, Roberto Calandra, Rowan McAllister, Sergey Levine. (2018) "Deep Reinforcement Learning in a Handful of Trials Using Probabilistic Dynamics Models." *NeurIPS 2018*. (Spotlight presentation, ~4% of submitted papers; NVIDIA Pioneer Award winner, 1/8) Presented at the Bay Area Machine Learning Symposium (Baylearn 2018) as a 15-minute talk.

Activities

Student Instructor

Fall 2023, Spring 2024

Introduction to Reinforcement Learning (Princeton Spring '24), Reasoning about Computation (Princeton Fall '23), Probability and Random Processes (Berkeley Fall '18 - Spring '19)

Planned/taught discussion sections, held office hours, and created/graded homeworks and exams.

Reviewer

NeurIPS 2021; International Conference on Learning Representations (ICLR) 2021; International Conference on Robotics and Automation (ICRA) 2021

Skills

ML Frameworks

Tensorflow, PyTorch, JAX

Languages

Python, Java, C/C++