

Michael Luo

☎ (408)-888-8377 | ✉ michael.luo@berkeley.edu | 🌐 www.michaelzhiluo.com | 🌐 [michaelzhiluo](https://github.com/michaelzhiluo) | in [michaelzhiluo](https://www.linkedin.com/in/michaelzhiluo)

EDUCATION

University of California: Berkeley Berkeley CA
PhD in Electrical Engineering and Computer Science, Advised by Ion Stoica Aug 2021 - Fall 2025

University of California: Berkeley Berkeley, CA
GPA: 3.98/4, M.S. in EECS + B.S. in EECS + Business Aug 2016 - May 2021

EXPERIENCE

Student Researcher & Part Time Sept 2023 - March 2025
Google Deepmind Mountain View, CA

- SaxML Team; Developed and researched efficient agentic systems for serving and post-training.

Software Development Engineer, Intern June 2020 - Jan 2021
Anyscale, Ray Core Team Berkeley, CA

- Scaled distributed RL training and developed asynchronous RL algorithms with Ray/RLlib.

Machine Learning Engineer, Intern June 2019 - Aug 2019
Amazon AI Palo Alto, CA

- Created the first embedding/RAG based recommendation system for Amazon Ads.

Computer Vision Engineer, Intern June 2018 - Aug 2018
Cisco Meraki, Smart Camera Team San Francisco, CA

- Developed model-based object detection to track individuals across Meraki cameras.

HONORS AND AWARDS

Neural Information Processing Systems Conference, Oral Award 2024
NeurIPS

- Awarded to *Stylus: Automatic Adapter Selection for Diffusion Models*.

USENIX Annual Technical Conference, Best Paper Award 2024
Usenix ATC

- Awarded to *Starburst: A Cost-aware Scheduler for Hybrid Cloud*

ACM International Collegiate Programming, Regional First Place 2016
Association for Computing Machinery

- Pacific Northwest Regional Division I

INVITED TALKS

DeepCoder: A Fully Open-Source 14B Coder at O3-mini Level 2025
Apple Cupertino, CA

DeepScaleR: Surpassing O1-Preview with a 1.5B Model by Scaling RL 2025
Google Deepmind, IBM Mountain View, CA

Autellix: An Efficient Serving Engine for LLM Agents as General Programs 2024
Meta Menlo Park, CA

Stylus: Automatic Adapter Selection 2024
SambaNova Palo Alto, CA

- [1] **Michael Luo**, Xiaoxiang Shi, Colin Cai, Tianjun Zhang, Justin Wong, Yichuan Wang, Chi Wang, Yanping Huang, Zhifeng Chen, Joseph E. Gonzalez, and Ion Stoica. “Autellix: An Efficient Serving Engine for LLM Agents as General Programs.” arXiv, 2025, arXiv:2502.13965.
- [2] **Michael Luo**, Sijun Tan, Roy Huang, Xiaoxiang Shi, Rachel Xin, Colin Cai, Ameen Patel, Alpay Ariyak, Qingyang Wu, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. DeepCoder: A Fully Open-Source 14B Coder at O3-mini Level. 2025, Notion Blog, <https://pretty-radio-b75.notion.site/DeepCoder-A-Fully-Open-Source-14B-Coder-at-O3-mini-Level-1cf81902c14680b3bee5eb349a512a51>.
- [3] **Michael Luo**, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. DeepScaleR: Surpassing O1-Preview with a 1.5B Model by Scaling RL. 2025, Notion Blog, <https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2>.
- [4] **Michael Luo**, Justin Wong, Brandon Trabucco, Yanping Huang, Joseph E. Gonzalez, Zhifeng Chen, Ruslan Salakhutdinov, and Ion Stoica. “Stylus: Automatic Adapter Selection.” *Advances in Neural Information Processing Systems*, vol. 37, 2024, pp. 32888–32915.
- [5] **Michael Luo**, Siyuan Zhuang, Suryaprakash Vengadesan, Romil Bhardwaj, Justin Chang, Eric Friedman, Scott Shenker, and Ion Stoica. “Starburst: A Cost-Aware Scheduler for Hybrid Cloud.” 2024 USENIX Annual Technical Conference (USENIX ATC 24), USENIX Association, July 2024, Santa Clara, CA, pp. 37–57.
- [6] **Michael Luo**, Jiahao Yao, Richard Liaw, Eric Liang, and Ion Stoica. “IMPACT: Importance Weighted Asynchronous Architectures with Clipped Target Networks.” 8th International Conference on Learning Representations (ICLR 2020).
- [7] **Michael Luo**, Ashwin Balakrishna, Brijen Thananjeyan, Suraj Nair, Julian Ibarz, Jie Tan, Chelsea Finn, Ion Stoica, and Ken Goldberg. MESA: Offline Meta-RL for Safe Adaptation and Fault Tolerance. *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [8] Zongheng Yang, Zhanghao Wu, **Michael Luo**, Wei-Lin Chiang, Romil Bhardwaj, Woosuk Kwon, Siyuan Zhuang, Frank Sifei Luan, Gautam Mittal, Scott Shenker, and Ion Stoica. “SkyPilot: An Intercloud Broker for Sky Computing.” 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23), USENIX Association, Apr. 2023, Boston, MA, pp. 437–455.
- [9] Eric Liang, Zhanghao Wu, **Michael Luo**, Sven Mika, Joseph E. Gonzalez, and Ion Stoica. “RLlib Flow: Distributed Reinforcement Learning Is a Dataflow Problem.” *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 5506–5517.
- [10] Ashwin Balakrishna, Brijen Thananjeyan, Suraj Nair, **Michael Luo**, Krishnan Srinivasan, Minho Hwang, Joseph Gonzalez, Julian Ibarz, Chelsea Finn, and Kenneth Goldberg. “Recovery RL: Safe Reinforcement Learning With Learned Recovery Zones.” *IEEE Robotics and Automation Letters*, vol. PP, 2021, pp. 1–1. <https://doi.org/10.1109/LRA.2021.3070252>
- [11] Justin Wong, Yury Orlovskiy, **Michael Luo**, Sanjit A. Seshia, and Joseph E. Gonzalez. SimpleStrat: Diversifying Language Model Generation with Stratification. arXiv, 2024. <https://arxiv.org/abs/2410.09038>.
- [12] Rachit Dubey, Erin Grant, Michael Luo, Karthik Narasimhan, and Thomas Griffiths. Connecting Context-specific Adaptation in Humans to Meta-learning. arXiv, 2020. <https://arxiv.org/abs/2011.13782>.
- [13] Ryan Hoque, Ashwin Balakrishna, Carl Putterman, **Michael Luo**, Daniel S. Brown, Daniel Seita, Brijen Thananjeyan, Ellen R. Novoseller, and Ken Goldberg. “LazyDagger: Reducing Context Switching in Interactive Imitation Learning.” *Proceedings of the 2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, IEEE, 2021, pp. 502–509. doi:10.1109/CASE49439.2021.9551469.

- [14] Xuanlin Li, Brandon Trabucco, Dong Huk Park, **Michael Luo**, Sheng Shen, Trevor Darrell, and Yang Gao. Discovering Non-monotonic Autoregressive Orderings with Variational Inference. International Conference on Learning Representations (ICLR), 2021. <https://openreview.net/forum?id=jP1vTH3inC>.
- [15] Zongheng Yang, Wei-Lin Chiang, Sifei Luan, Gautam Mittal, **Michael Luo**, and Ion Stoica. “Balsa: Learning a Query Optimizer Without Expert Demonstrations.” Proceedings of the 2022 International Conference on Management of Data (SIGMOD ’22), Association for Computing Machinery, 2022, pp. 931–944.
- [16] Jeffrey Ichnowski, Paras Jain, Bartolomeo Stellato, Goran Banjac, **Michael Luo**, Francesco Borrelli, Joseph E. Gonzalez, Ion Stoica, and Ken Goldberg. “Accelerating Quadratic Optimization with Reinforcement Learning.” Advances in Neural Information Processing Systems, vol. 34, 2021, pp. 21043–21055.

ACADEMIC SERVICES

Reviewer	2021-2025
<i>Neural Information Processing Systems Conference (NeurIPS)</i>	
Reviewer	2021-2023
<i>International Conference on Learning Representations (ICLR)</i>	

TECHNICAL SKILLS

Areas: Deep Reinforcement Learning (RL), Post-Training, LLM Serving, Distributed Systems, Natural Language Processing, Robotics, Computer Vision

Languages: Python, C/C++, CUDA, Java, SQL (Postgres), Javascript, C#, Golang, Bash, HTML/CSS

ML Libraries: PyTorch, CUDA, JAX, Ray, Tensorflow, XLA, VeRL

Infrastructure: Kubernetes, Docker, Terraform, Spark, Hadoop, Google Cloud Platform (GCP), AWS, Azure