Bingbin Liu | Curriculum Vitae

bliu@g.harvard.edu • clarabing.github.io

Employment

- **Kempner Institute at Harvard University** *2025 – present*
  Postdoc Research Fellow.
- **Simons Institute at UC Berkeley** *Fall 2024*
  Postdoc Research Fellow at Special Year on LLMs and Transformers, Part 1 and Modern Paradigms in Generalization.

Education

- **Carnegie Mellon University** *2019 – 2024*
  Ph.D. in Machine Learning advised by Andrej Risteski and Pradeep Ravikumar.
- **Stanford University** *2017 – 2019*
  M.S. in Computer Science advised by Fei-Fei Li.
- **The University of Hong Kong** *2013 – 2017*
  B.Eng. Major in Computer Science and Minor in Math. First-Class Honour.

Publications and Preprints

- **In Good GRACEs: Principled Teacher Selection for Knowledge Distillation**
  In submission [link]
  Abhishek Panigraphi, Bingbin Liu, Sadhika Malladi, Sham M. Kakade, Surbhi Goel
- **Adam or Gauss-Newton? — A Comparative Study In Terms of Basis Alignment and SGD Noise**
  In submission [link]
  Bingbin Liu, Rachit Bansal, Depen Morwani, Nikhil Vyas, David Alvarez-Melis, Sham M. Kakade
- **Progressive distillation induces an implicit curriculum**
  *ICLR 2025 (**Oral**)* [link]
  Abhishek Panigrahi*, Bingbin Liu*, Sadhika Malladi, Andrej Risteski, Surbhi Goel
- **Understanding Augmentation-based Self-Supervised Representation Learning via RKHS Approximation**
  *ICLR 2024 (**Spotlight**)* [link]
  Runtian Zhai, Bingbin Liu, Andrej Risteski, Zico Kolter, Pradeep Ravikumar
- **TinyGSM: achieving $>80\%$ on GSM8k with small language models**
  preprint [link]
  Bingbin Liu, Sebastien Bubeck, Ronen Eldan, Janardhan Kulkarni, Yuanzhi Li, Anh Nguyen, Rachel Ward, Yi Zhang
- **Exposing Attention Glitches with Flip-Flop Language Modeling**
  *NeurIPS 2023 (**Spotlight**)* [link]
  Bingbin Liu, Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, Cyril Zhang
- **Transformers are uninterpretable with myopic methods: a case study with bounded Dyck grammars**
  NeurIPS 2023 [link]
  Kaiyue Wen, Yuchen Li, Bingbin Liu, Andrej Risteski
- **Transformers Learn Shortcuts to Automata**
  *ICLR 2023 (**Oral**)* [link]
  Bingbin Liu, Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, Cyril Zhang
- **Masked prediction tasks: a parameter identifiability view**
  NeurIPS 2022 [link]
  Bingbin Liu, Daniel Hsu, Pradeep Ravikumar, Andrej Risteski

- **Analyzing and Improving the Optimization Landscape of Noise-Contrastive Estimation**
  *ICLR 2022 (**Spotlight**)* [link]
  Bingbin Liu, Elan Rosenfeld, Pradeep Ravikumar, Andrej Risteski
- **Contrastive learning of strong-mixing continuous-time stochastic processes**
  AISTATS 2021; abstract version presented at the WiML workshop colocated with NeurIPS 2020.[link]
  Bingbin Liu, Pradeep Ravikumar, Andrej Risteski
- **Generalized Boosting**
  NeurIPS 2020 [link]
  Arun Sai Suggala, Bingbin Liu, Pradeep Ravikumar
- **Spatiotemporal Relationship Reasoning for Pedestrian Intent Prediction**
  ICRA 2020, IEEE-RAL [link].
  Bingbin Liu, Ehsan Adeli, Zhangjie Cao, Kuan-Hui Lee, Abhijeet Shenoi, Adrien Gaidon, Juan Carlos Niebles
- **A Computer Vision System for Deep Learning-based Detection of Patient Mobilization Activities in the ICU**
  Nature Digital Medicine 2019 [link]
  Serena Yeung*, Francesca Rinaldo*, Jeffrey Jopling, Bingbin Liu, Rishab Mehra, Lance Downing, Michelle Guo, Gabriel Bianconi, Alexandre Alahi, Julia Lee, Brandi Campbell, Kayla Deru, William Beninati, Li Fei-Fei, Arnold Milstein
- **Learning to Decompose and Disentangle Representations for Video Prediction**
  NeurIPS 2018 [link]
  Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li Fei-Fei, Juan Carlos Niebles
- **Temporal Modular Networks for Retrieving Complex Compositional Activities in Videos**
  ECCV 2018 [link]. Also presented at WiCV 2018.
  Bingbin Liu, Serena Yeung, Edward Chou, De-An Huang, Li Fei-Fei, Juan Carlos Niebles
- **3D Point Cloud-Based Visual Prediction of ICU Mobility Care Activities**
  MLHC 2018 [link]
  Bingbin Liu*, Michelle Guo*, Edward Chou, Rishab Mehra, Serena Yeung, N. Lance Downing, Francesca Rinaldo, Jeffrey Jopling, Branci Campbell, Kayla Deru, William Beninati, Arnold Milstein, Li Fei-Fei
- **Descriptive Analysis of ICU Patient Mobilization from Depth Videos**
  ML4H 2018 (colocated with NeurIPS 2018)
  Laëtitia Shao*, Zaid Nabulsi*, Ruchir Rastogi*, Bingbin Liu, Francesca Rinaldo, Serena Yeung, N. Lance Downing, William Beninati, Arnold Milstein, Li Fei-Fei

Research Experience

- **Microsoft Research Lab – Redmond** *Summer 2023*
  Mentored by Yi Zhang, Yuanzhi Li, Rachel Ward, Ronen Eldan, Sébastien Bubeck.
- **Microsoft Research Lab – New York City** *Summer 2022*
  Mentored by Cyril Zhang, Akshay Krishnamurthy, Surbhi Goel, Jordan T. Ash.
  Transformers for algorithmic reasoning tasks defined with finite-state automata.
- **Stanford Vision and Learning Lab** *Stanford University, 2017 – 2019*
  Advised by Fei-Fei Li and Juan Carlos Niebles.
  Compositional video understanding with graphs, temporal modules, and disentangled representations.
- **Partnership in AI Assisted Care (PAC)** *Stanford University, 2017 – 2019*
  Advised by Fei-Fei Li, Arnold Milstein, Serena Yeung.
  Student lead of the ICU team: using depth videos to monitor and document patient mobilization levels in ICU care.
- **HKU Computer Vision Lab** *The University of Hong Kong, 2016 – 2017*
  Advised by Kenneth K.Y. Wong.
  Final year project on improving video object recognition with temporal context. Summer project on cell segmentation.
- **UCSB Programming Language Lab** *UC Santa Barbara, Feb – June 2016*

Advised by Ben Hardekopf.
Undergraduate study: building a compiler in Haskell of a functional ISA for simplified formal verifications.

Other Experiences

- **Organizing**
  - LeT-All Fall 2025 Mentorship Workshop.
  - Methods and Opportunities at Small Scale (MOSS) at ICML 2025.
  - Tutorial "Sandbox for the Blackbox: How LLMs Learn Structured Data" at NeurIPS 2024.
  - Mathematics of Modern Machine Learning (M3L) workshop at NeurIPS 2023, 2024.
- **Reviewing**
  Reviewer for ICLR, ICML, NeurIPS (top 8% reviewer at NeurIPS 2022), COLT, AISTATS.
- **Graduate teaching assistant** *Carnegie Mellon University*
  - 10716 - Advanced Machine Learning, Spring 2022.
  - 10707 - Advanced Deep Learning, Spring 2021.
- **Graduate teaching assistant** *Stanford University*
  - CS231N - Convolutional Neural Networks for Visual Recognition, Spring 2018, Spring 2019.
  - MED277/CS337 - AI-Assisted Health Care, Fall 2018, Winter 2019.
- **Mentoring and outreach**
  - Co-organizer of Women+ mentoring group at Simons Institute, Fall 2024.
  - CMU AI Undergraduate Mentoring Program at Carnegie Mellon University, 2020, 2023.
  - AI4All at Stanford University, Summer 2018.
  - Girls teach Girls to Code at Stanford University, Spring 2018.

Awards

- **Tang Family Presidential Fellowship** *2023*
- **Powering a Sustainable Generation Scholarship by CLP** *2015*
- **HKU World Wide Scholarship** *2015*
- **Dean's Honours List** *2013 - 2017*
- **Lee Shau Kee Scholarships** *2013 - 2017*