Yuchen Li

Email: yuchenl4@cs.cmu.edu **Website:** https://www.cs.cmu.edu/~yuchenl4/
LinkedIn: https://www.linkedin.com/in/yuchenli01/ **Twitter:** https://twitter.com/_Yuchen_Li_
Work eligibility and visa status: legally authorized to work in the United States without sponsorship

EDUCATION

Carnegie Mellon University August 2020 – December 2025 (expected)
- Ph.D. student in Machine Learning, advised by Prof. Andrej Risteski Pittsburgh, PA
- Research interests: machine learning, natural language processing, and data mining GPA: 4.24/4.33
- I work on (1) improving mathematical understanding of language models (training dynamics, efficient sampling, mechanistic interpretability), and (2) developing principled approaches to self-supervised learning

Simons Institute for the Theory of Computing August 2024 - December 2024
- Visiting graduate student for the "Modern Paradigms in Generalization" program Berkeley, CA

University of Illinois at Urbana-Champaign August 2015 - May 2019
- B.S. Statistics and Computer Science; second major in Mathematics; minor in Electrical Engineering Urbana-Champaign, IL
- Research mentors: Prof. Jiawei Han, Prof. AJ Hildebrand, Prof. Pramod Viswanath GPA: 4.0/4.0
- Summa Cum Laude; Highest Distinction in both majors
- Bronze Tablet Scholar: top 3% in the graduating class
- John R. Pasta Outstanding Undergraduate Award 2018: 2 out of about 400 computer science undergraduate junior students

WORK

***Research Intern,* Microsoft Corporation** May 2024 - August 2024
- Working on understanding "data quality" for language model post training New York City, NY, United States
- Identified that the inclusion of "soft" data pairs is key for enabling contrastive preference optimization to outperform behavior cloning

***Student Researcher Intern,* Google LLC** May 2023 - December 2023
- Worked on efficient parallel decoding / sampling for language models Remote and New York City, NY, United States
- Proposed a theoretical framework for reasoning about generative masked language models, and empirical guidelines for training them

***Research Intern,* ByteDance Ltd. (parent company of TikTok)** May 2022 - August 2022
- Worked on a research project about machine translation at the ByteDance AI Lab Remote, United States
- Investigating the effect of language similarity in machine learning models for multi-lingual translation

***Machine Learning Engineer,* Quora, Inc.** August 2019 - August 2020
- Developed neural network models for predicting ads clickthrough rate Mountain View, CA, United States
- Led team-wide initiative of improving system robustness and reducing on-call burden in the Ads Ranking team
- Developed features for modeling the ads and the users
- Improved backend for cost-efficient low-latency online serving of machine learning models
- Launched online A/B testing to compare different model variants based on a set of key business metrics

***Software Engineer Intern,* Facebook, Inc.** May 2018 - August 2018
- Designed and implemented a platform for investigating machine learning feature importance Menlo Park, CA
- Proposed and developed feature perturbation and statistical analysis methods

PUBLICATIONS

(* indicates equal contribution or alphabetical order)

1. Dhruv Rohatgi, Abhishek Shetty, Donya Saless, **Yuchen Li**, Ankur Moitra, Andrej Risteski, Dylan J. Foster. *Taming Imperfect Process Verifiers: A Sampling Perspective on Backtracking*. arXiv 2510.03149, 2025.
2. Edoardo Botta*, **Yuchen Li***, Aashay Mehta*, Jordan T. Ash, Cyril Zhang, Andrej Risteski. *On the Query Complexity of Verifier-Assisted Language Generation*. International Conference on Machine Learning (ICML) 2025.
3. **Yuchen Li**, Alexandre Kirchmeyer*, Aashay Mehta*, Yilong Qin*, Boris Dadachev*, Kishore Papineni*, Sanjiv Kumar, Andrej Risteski. *Promises and Pitfalls of Generative Masked Language Modeling: Theoretical Framework and Practical Guidelines*. International Conference on Machine Learning (ICML) 2024.
4. Kaiyue Wen, **Yuchen Li,** Bingbin Liu, Andrej Risteski. *Transformers are uninterpretable with myopic methods: a case study with bounded Dyck grammars*. Conference on Neural Information Processing Systems (NeurIPS) 2023.
5. **Yuchen Li,** Yuanzhi Li, Andrej Risteski. *How Do Transformers Learn Topic Structure: Towards a Mechanistic Understanding*. International Conference on Machine Learning (ICML) 2023.
6. Ashwini Pokle *, Jinjin Tian *, **Yuchen Li** *, Andrej Risteski. *Contrasting the landscape of contrastive and non-contrastive learning*. Conference on Artificial Intelligence and Statistics (AISTATS) 2022.
7. **Yuchen Li,** Andrej Risteski. *The Limitations of Limited Context for Constituency Parsing*. Association for Computational Linguistics (ACL) 2021.
8. Xinwei He *, A.J. Hildebrand *, **Yuchen Li** *, Yunyi Zhang *. *Complexity of Leading Digit Sequences*. Journal of Discrete Mathematics & Theoretical Computer Science, vol. 22 no. 1, Automata, Logic and Semantics 2020.
9. Yu Shi *, Jiaming Shen *, **Yuchen Li**, Naijing Zhang, Xinwei He, Zhengzhi Lou, Qi Zhu, Matthew Walker, Myunghwan Kim and Jiawei Han. *Discovering Hypernymy in Text-Rich Heterogeneous Information Network by Exploiting Context Granularity*. Conference on Information and Knowledge Management (CIKM) 2019.
10. Hongyu Gong, **Yuchen Li**, Suma Bhat and Pramod Viswanath. *Context-Sensitive Malicious Spelling Error Correction*. The Web Conference (WWW) 2019.
11. **Yuchen Li** *, Zhengzhi Lou *, Yu Shi, and Jiawei Han. *Temporal Motifs in Heterogeneous Information Networks*. In Proceedings of the 14th International Workshop on Mining and Learning with Graphs (MLG) 2018.

TALKS

Towards Mathematical Understanding of Modern Language Models

- CMU Artificial Intelligence Seminar Series — Pittsburgh, PA, February 2025
- Flatiron Institute Transformers Reading Group — New York City, NY, August 2024
- Seminar on Foundational Artificial Intelligence — Virtual, April 2024

Transformers are uninterpretable with myopic methods: a case study with bounded Dyck grammars

- DeepMath Conference on the Mathematical Theory of Deep Neural Networks — Virtual, November 2023
  - Recording: https://www.youtube.com/live/BZ34XH4ffEc?si=LjXAiYywQWys4zG5&t=20360

How Do Transformers Learn Topic Structure: Towards a Mechanistic Understanding

- Prof. Graham Neubig's lab (NeuLab), CMU Language Technologies Institute — Pittsburgh, PA, May 2023
- Machine Learning Reading Group, University of Illinois at Urbana-Champaign (UIUC) — Virtual, April 2023

The Limitations of Limited Context for Constituency Parsing

- Association for Computational Linguistics (ACL) Conference — Virtual, August 2021
- NEC Laboratories Europe — Virtual, July 2021
- Approximately Correct Machine Intelligence (ACMI) Lab, CMU — Virtual, June 2021

TEACHING

***Teaching Assistant,* Advanced Introduction to Machine Learning** August 2022 - December 2022
- Taught one recitation session on optimization Pittsburgh, PA
- Held weekly office hours; created and graded homework problems
- Course webpage: https://www.cs.cmu.edu/~nihars/teaching/10715-Fa22/index.html
- Instructor: Prof. Nihar Shah

***Teaching Assistant,* Probabilistic Graphical Models** January 2022 - May 2022
- Taught two recitation sessions (about algorithmic reduction and causality, respectively) Pittsburgh, PA
- Mentored nine teams of student projects (topics include natural language processing, information retrieval, 3D computer vision, and causality)
- Held weekly office hours; created and graded homework and quiz problems
- Course webpage: https://andrejristeski.github.io/10708-22/
- Instructors: Prof. Andrej Risteski and Prof. Hoda Heidari

***Volunteer Teaching Assistant,* Data Science Bootcamp** May 2019
- Contributed to creating a Python tutorial focusing on NumPy and Matplotlib libraries. Urbana, IL
- Helped in the class section for about 20 math graduate students and answered their programming questions
- Program webpage: https://ravat1.github.io/2019DSB-website/
- Instructor: Prof. Uma Ravat

***Member,* Eta Kappa Nu (HKN) Honors Society Alpha Chapter at UIUC** August 2016 - May 2019
- Contributed to its educational website Weber's Wiki Urbana, IL
- Developed and taught two Analog Signal Processing midterm review sessions, each with about 40 students attending

SERVICE

***Reviewer,* Academic Research Conferences**
- International Conference on Machine Learning (ICML) 2024, 2025
- Journal of Machine Learning Research (JMLR) 2024, 2025
- Artificial Intelligence and Statistics (AISTATS) 2022, 2024
- International Conference on Learning Representations (ICLR) 2023, 2024
- Conference on Neural Information Processing Systems (NeurIPS) 2021, 2023
- Mathematics of Modern Machine Learning Workshop at NeurIPS 2023
- Knowledge and Logical Reasoning Workshop at ICML 2023

***Student Member,* CMU Machine Learning Peers** 2025
- Mentored junior PhD students in CMU Machine Learning Department on topics including finding an advisor, maintaining relationship with advisor and collaborators, course selection, experience at CMU and in Pittsburgh, and mental health Pittsburgh, PA

***Volunteer,* Conference on Neural Information Processing Systems (NeurIPS)** 2023
- Checked in conference attendees and speakers at the registration desk New Orleans, LA
- Assisted with poster display and other organizational logistics at a workshop room

***Session Chair,* Learning Theory Alliance Mentorship Workshop** 2023
- Volunteered in this mentorship workshop centered on communicating one's research verbally for undergraduate, masters, PhD students, and postdocs interested in learning theory and adjacent fields Online

- Hosted a session in which participants gave research talks and the audience discussed learnings and provided feedback

***Mentor,* CMU Computer Science Graduate Application Support Program** 2020
- Connected with potential applicants, particularly those from underrepresented groups Pittsburgh, PA
- Offered feedbacks on their application materials and answered their questions related to graduate school application

***Student Member,* CMU Machine Learning Graduate Programs Admissions Committee** multiple years
- Assisted in faculty-led tasks and meetings for reviewing and admitting applicants Pittsburgh, PA
- Contributed preliminary reviews for over 100 applicants to CMU Machine Learning Masters and PhD programs

HONORS

- Citadel / Correlation One Terminal Al programming contest midwest regional: 3rd place ($2500 award for team) State level, 2022
- Conference Travel Grant: for top undergraduate students presenting their research projects University level, 2018
- University of Illinois Undergraduate Math Contest: 4th place University level, 2016
- Virginia Tech Regional Math Contest: ranked 17 out of 701 participants all over the USA National level, 2015
- American Invitational Mathematics Examination II (AIME II): 11th Grade 97.59% worldwide International level, 2014
- American Mathematics Competitions 10B (AMC 10B): the 11th place out of 37394 participants globally International level, 2013

SKILLS

- Proficient in *Python, PyTorch, TensorFlow, Jax, C, SQL, MATLAB, PHP*
- Intermediate in *C++, R, Haskell*
- Basic in *Java, JavaScript*