

Alisa Liu

Ph.D. Candidate, University of Washington
alisaliu@cs.washington.edu | [alisawuffles.github.io](https://github.com/alisawuffles) | 

Education

2020 – *present* **Ph.D. candidate** in Computer Science, University of Washington
Advisors: Yejin Choi, Noah A. Smith

2020 **B.A.**, Northwestern University
Double Major: Computer Science and Math
GPA: 4.0/4.0, *summa cum laude*

Employment

Dec 2024 – *present* Academic Collaborator, NVIDIA
Host: Yejin Choi

Jun – Dec 2023 Student Researcher Intern, Google
Host: Steven Zheng, Swaroop Mishra

Jun – Sep 2020 Undergraduate Research Intern, Allen Institute for AI
Host: Chandra Bhagavatula

Awards & Honors

2025 Rising Stars in Data Science @ Stanford

2024 OpenAI SuperAlignment Fast Grant (\$150,000)

2023 Inverse Scaling Challenge Third Prize for “Memorization Traps” (\$5,000)

2022 NSF Graduate Research Fellowship (GRFP)

2020 Hacherl Endowed Graduate Fellowship granted by Paul G. Allen School

2020 Northwestern Outstanding Computer Science Senior

2020 CRA Outstanding Undergraduate Researcher Award Honorable Mention

Selected Publications

* denotes equal contribution

Tokenization for Language Models

A Liu*, J Hayase*, V Hofmann, S Oh, N Smith, Y Choi. [SuperBPE: Space Travel for Language Models](#). *COLM* (2025).

B Zheng, A Liu, O Ahia, J Hayase, Y Choi, N Smith. [Broken Tokens? Your Language Model can Secretly Handle Non-Canonical Tokenizations](#). *NeurIPS (Spotlight)* (2025).

H Xu, A Liu, J Hayase, Y Choi, N Smith. [Are you going to finish that? A Practical Study of the Partial Token Problem](#). 2025.

J Hayase, A Liu, N Smith, S Oh. [Sampling from Your Language Model One Byte at a Time](#). 2025.

A Liu*, J Hayase*, Y Choi, S Oh, N Smith. [Data Mixture Inference: What do BPE Tokenizers Reveal about their Training Data?](#) *NeurIPS* (2024).

Text Generation & Data Creation

A Liu, X Han, Y Wang, Y Tsvetkov, Y Choi, N Smith. [Tuning Language Models by Proxy](#). *COLM (Spotlight)* (2024).

R Shi, Y Chen, Y Hu, A Liu, H Hajishirzi, N Smith, S Du. [Decoding-Time Language Model Alignment with Multiple Objectives](#). *NeurIPS* (2024).

M Zhang, O Press, W Merrill, A Liu, N Smith. [How Language Model Hallucinations Can Snowball](#). *ICML* (2024).

A Liu, Z Wu, J Michael, A Suhr, P West, A Koller, S Swayamdipta, N Smith, Y Choi. [We're Afraid Language Models Aren't Modeling Ambiguity](#). *EMNLP* (2023).

Y Wang, Y Kordi, S Mishra, A Liu, N Smith, D Khashabi, H Hajishirzi. [Self-Instruct: Aligning Language Models with Self-Generated Instructions](#). *ACL* (2023).

J Lee, A Liu, O Ahia, H Gonen, N Smith. [That was the last straw, we need more: Are Translation Systems Sensitive to Disambiguating Context?](#) *Findings of EMNLP* (2023).

A Liu, S Swayamdipta, N Smith, Y Choi. [WANLI: Worker and AI Collaboration for Natural Language Inference Dataset Creation](#). *Findings of EMNLP* (2022).

J Liu, A Liu, X Lu, S Welleck, P West, R Le Bras, Y Choi, H Hajishirzi. [Generated Knowledge Prompting for Commonsense Reasoning](#). *ACL* (2022).

A Liu, M Sap, X Lu, S Swayamdipta, C Bhagavatula, N Smith, Y Choi. [DExperts: Decoding-Time Controlled Text Generation with Experts and Anti-Experts](#). *ACL* (2021).

Language Model Development

NVIDIA, ..., A Liu, [NVIDIA Nemotron 3: Efficient and Open Intelligence](#). 2025.

Olmo Team, ..., A Liu, ..., A Farhadi, N Smith, H Hajishirzi. [Olmo 3](#). 2025.

N Lambert, ..., A Liu, ..., H Hajishirzi. [TüLU 3: Pushing Frontiers in Open Language Model Post-Training](#). *COLM* (2025).

I McKenzie, ..., A Liu, ..., S Bowman, E Perez. [Inverse Scaling: When Bigger Isn't Better](#). *TMLR (Featured)* (2023).

Service, Outreach & Teaching

Invited Talks Panelist @ Tokenization Workshop @ ICML 2025, LLM Merging Competition @ NeurIPS 2024, Cohere for AI, Microsoft Turing, Microsoft Translator, Square, UVA, UCLA, UMass Amherst, UBC, Northeastern, Northwestern, Fondazione Bruno Kessler, MIT

Teaching Instructor of NLP Capstone (2025) @ UW, TA for NLP (2024) and NLP Seminar (2023) @ UW, TA for Algorithms (2019) and Machine Learning (2019) @ Northwestern

Organization Tutorial on Synthetic Data Generation in the Era of LLMs @ ACL 2025, Workshop on Understanding Implicit and Underspecified Language @ EACL 2024

Department Service Co-Chair of DEIA Committee (2024 – 2025), DEIA Special Considerations Reader for Ph.D. program (2025), Pre-Application Mentoring Service (2020 – 2024), Ph.D. Application Reader (2020 – 2024), Graduate Student Volunteer Coordinator (2023 – 2024), Co-Chair of Care Committee (2021 – 2024), Co-Organizer of New Grad Orientation (2021)

Mentorship Brian Zheng (UW), Hao Xu (UW), David Pohl (Institute of Science Tokyo), Andrew Shaw (UW), Margaret Li (UW), Skyler Hallinan (UW → USC), Muru Zhang (UW → USC)