

Tianyu Gao

tianyug@princeton.edu | gaotianyu.xyz/about

Education

Princeton University

Princeton, New Jersey, USA

Ph.D. Candidate in Computer Science

2020 - 2025 (expected)

- Research area: natural language processing, large language models, machine learning
- Advisor: [Prof. Danqi Chen](#)

Tsinghua University

Beijing, China

B.Eng. in Computer Science and Technology

2016 - 2020

- Advisor: [Prof. Zhiyuan Liu](#)

Honors and Awards

- 2023 IBM PhD Fellowship (10 recipients globally)
- 2023 The School of Engineering and Applied Science (SEAS) Award for Excellence, Princeton University (four recipients from computer science)
- 2023 Qualcomm Innovation Fellowship Finalist
- 2022 ACL Outstanding Paper Award (top 1.1% out of 1,062 accepted papers)
- 2019 Special-Class Scholarship, Tsinghua University (10 recipients university-wide; one recipient from computer science)
- 2019 Cai Xiong Scholarship, Tsinghua University (10 recipients university-wide, two recipients from computer science)
- 2019 Fellowship of the Initiative Scientific Research Program, Tsinghua University (10 recipients university-wide; two recipients from computer science)

Publications

** indicates equal contribution. "The most cited papers" are based on Semantic Scholar results in 2024/11.*

[\[Google Scholar\]](#) [\[Semantic Scholar\]](#)

Preprints

- [1] *Metadata Conditioning Accelerates Language Model Pre-training.*
Tianyu Gao, Alexander Wettig, Luxi He, Yihe Dong, Sadhika Malladi, Danqi Chen. 2025.
- [2] *LongProc: Benchmarking Long-Context Language Models on Long Procedural Generation.*
Xi Ye, Fangcong Yin, Yinghui He, Joie Zhang, Howard Yen, **Tianyu Gao**, Greg Durrett, Danqi Chen. 2025.
- [3] *How to Train Long-Context Language Models (Effectively).*
Tianyu Gao*, Alexander Wettig*, Howard Yen, Danqi Chen. 2024.
- [4] *HELMET: How to Evaluate Long-Context Language Models Effectively and Thoroughly.*
Howard Yen, **Tianyu Gao**, Minmin Hou, Ke Ding, Daniel Fleischer, Peter Izsak, Moshe Wasserblat, Danqi Chen. 2024.
- [5] *Improving Language Understanding from Screenshots.*
Tianyu Gao, Zirui Wang, Adithya Bhaskar, Danqi Chen. 2024.

Peer-reviewed Conference/Journal Papers

- [1] *LitSearch: A Retrieval Benchmark for Scientific Literature Search.*
Anirudh Ajith, Mengzhou Xia, Alexis Chevalier, Tanya Goyal, Danqi Chen, **Tianyu Gao**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [2] *Long-Context Language Modeling with Parallel Context Encoding.*
Howard Yen, **Tianyu Gao**, Danqi Chen. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [3] *Sheared LLaMA: Accelerating Language Model Pre-training via Structured Pruning.*
Mengzhou Xia, **Tianyu Gao**, Zhiyuan Zeng, Danqi Chen. In *Proceedings of The International Conference on Learning Representations (ICLR)*, 2024.
- [4] *Evaluating Large Language Models at Evaluating Instruction Following.*
Zhiyuan Zeng, Jiatong Yu, **Tianyu Gao**, Yu Meng, Tanya Goyal, Danqi Chen. In *Proceedings of The International Conference on Learning Representations (ICLR)*, 2024.
- [5] *Enabling Large Language Models to Generate Text with Citations.*
Tianyu Gao, Howard Yen, Jiatong Yu, Danqi Chen. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- [6] *Fine-Tuning Language Models with Just Forward Passes.*
Sadhika Malladi*, **Tianyu Gao***, Eshaan Nichani, Alex Damian, Jason D. Lee, Danqi Chen, Sanjeev Arora. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2023. **Oral presentation (top 2% out of 3,218 accepted papers).**
- [7] *What In-Context Learning "Learns" In-Context: Disentangling Task Recognition and Task Learning.*
Jane Pan, **Tianyu Gao**, Howard Chen, Danqi Chen. In *Findings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [8] *The CRINGE Loss: Learning What Language Not to Model.*
Leonard Adolphs, **Tianyu Gao**, Jing Xu, Kurt Shuster, Sainbayar Sukhbaatar, Jason Weston. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [9] *Should You Mask 15% in Masked Language Modeling.*
Alexander Wettig*, **Tianyu Gao***, Zexuan Zhong, Danqi Chen. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, 2023. **EACL 2023's 2nd most cited paper.**
- [10] *Ditch the Gold Standard: Re-evaluating Conversational Question Answering.*
Huihan Li*, **Tianyu Gao***, Manan Goenka, Danqi Chen. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022. **Outstanding Paper Award (top 1.1% out of 1,062 accepted papers).**
- [11] *Recovering Private Text in Federated Learning of Language Models.*
Samyak Gupta*, Yangsibo Huang*, Zexuan Zhong, **Tianyu Gao**, Kai Li, Danqi Chen. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- [12] *Automatic Label Sequence Generation for Prompting Sequence-to-sequence Models.*
Zichun Yu, **Tianyu Gao**, Zhengyan Zhang, Yankai Lin, Zhiyuan Liu, Maosong Sun, Jie Zhou. In *Proceedings of the International Conference on Computational Linguistics (COLING)*, 2022.
- [13] *SimCSE: Simple Contrastive Learning of Sentence Embeddings.*
Tianyu Gao*, Xingcheng Yao*, Danqi Chen. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021. **EMNLP 2021's 2nd most cited paper.**
- [14] *Making Pre-trained Language Models Better Few-shot Learners.*
Tianyu Gao*, Adam Fisch*, Danqi Chen. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, 2021. **ACL 2021's 2nd most cited paper.**
- [15] *Manual Evaluation Matters: Reviewing Test Protocols of Distantly Supervised Relation Extraction.*
Tianyu Gao, Xu Han, Keyue Qiu, Yuzhuo Bai, Zhiyu Xie, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou. In *Findings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.

- [16] *KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation.*
Xiaozhi Wang, **Tianyu Gao**, Zhaocheng Zhu, Zhiyuan Liu, Juanzi Li, Jian Tang. *Transactions of the Association for Computational Linguistics*, 2021.
- [17] *Learning from Context or Names? An Empirical Study on Neural Relation Extraction.*
Hao Peng*, **Tianyu Gao***, Xu Han, Yankai Lin, Peng Li, Zhiyuan Liu, Maosong Sun, Jie Zhou. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [18] *More Data, More Relations, More Context and More Openness: A Review and Outlook for Relation Extraction.*
Xu Han*, **Tianyu Gao***, Yankai Lin*, Hao Peng, Yaoliang Yang, Chaojun Xiao, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou. In *Proceedings of the Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, 2020.
- [19] *Continual Relation Learning via Episodic Memory Activation and Reconsolidation.*
Xu Han, Yi Dai, **Tianyu Gao**, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2020.
- [20] *Few-shot Relation Extraction via Bayesian Meta-learning on Task Graphs.*
Meng Qu, **Tianyu Gao**, Louis-Pascal Xhonneux, Jian Tang. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.
- [21] *Neural Snowball for Few-Shot Relation Learning.*
Tianyu Gao, Xu Han, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, Maosong Sun. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [22] *FewRel 2.0: Towards More Challenging Few-Shot Relation Classification.*
Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, Jie Zhou. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [23] *OpenNRE: An Open and Extensible Toolkit for Neural Relation Extraction.*
Xu Han*, **Tianyu Gao***, Yuan Yao, Deming Ye, Zhiyuan Liu, Maosong Sun. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, 2019.
- [24] *Hybrid Attention-Based Prototypical Networks for Noisy Few-Shot Relation Classification.*
Tianyu Gao*, Xu Han*, Zhiyuan Liu, Maosong Sun. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

Research and Industry Experience

Facebook AI Research (FAIR) <i>Research Intern (Mentors: Jason Weston and Mike Lewis)</i>	Jun. 2022 – Aug. 2022 New York City, NY, USA
Mila-Quebec AI Institute <i>Research Intern (Mentor: Prof. Jian Tang)</i>	Jul. 2019 – Sept. 2019 Montreal, Canada
WeChat AI, Tencent <i>Tencent Rhino-Bird Elite Training Program</i>	May 2019 - May 2020 Beijing, China
Momenta <i>Research Intern</i>	May 2017 - May 2018 Beijing, China

Teaching Experience

COS484: Natural Language Processing <i>Teaching assistant. Instructor: Prof. Karthik Narasimhan</i>	Spring 2022 Princeton University
COS324: Introduction to Machine Learning <i>Teaching assistant. Instructors: Prof. Sanjeev Arora and Prof. Danqi Chen</i>	Fall 2021 Princeton University

Invited Talks/Lectures

- 2024/11 **MIT/University of Waterloo/University of Queensland.**
Title: Developing Effective Long-Context Language Models.
- 2024/10 **Princeton Research Computing Workshop.**
Title: Using Large Language Models in your Research: Fine-Tuning, Embeddings, and Using APIs.
- 2024/10 COS 597R (Deep Dive into Large Language Models; instructor: Prof. Danqi Chen and Prof. Sanjeev Arora), **Princeton University.**
Title: Long-Context Language Models.
- 2024/09 **Google Deepmind.**
Title: A Recipe for Effectively Extending Context Windows for Language Models.
- 2024/04 CS 6501 (Natural Language Processing; instructor: Prof. Yu Meng), **UVA.**
Title: Language Models as the New Search Engines.
- 2024/04 CSCI-SHU 376 (Natural Language Processing; instructor: Prof. Chen Zhao), **NYU Shanghai.**
Title: Language Models as the New Search Engines.
- 2024/03 15-442/642 (Machine Learning Systems; instructor: Prof. Zhihao Jia), **CMU.**
Title: Language Models as the New Search Engines.
- 2024/03 **FAI-seminar.**
Title: Long-Context Language Modeling with Parallel Context Encoding.
- 2024/02 **UMass NLP Seminar.**
Title: Enabling Large Language Models to Generate Text with Citations.
- 2024/02 COS 598D (Advanced Topics in Computer Science: Systems and Machine Learning; instructor: Prof. Kai Li), **Princeton University.**
Title: Pre-trained Large Language Models.
- 2023/11 **Tsinghua NLP Seminar.**
Title: Enabling LLMs to Generate Text with Citations.
- 2023/06 **Semantic Machines.**
Title: Enabling Large Language Models to Generate Text with Citations.
- 2023/06 **Apple.**
Title: Enabling Large Language Models to Generate Text with Citations.
- 2023/06 **SingleStore.**
Title: Enabling Large Language Models to Generate Text with Citations.
- 2023/06 **MLNLP Seminar.**
Title: Enabling Large Language Models to Generate Text with Citations (in Chinese).
- 2022/06 **AI Time.**
Title: Ditch the Gold Standard: Re-evaluating Conversational Question Answering (in Chinese).
- 2021/07 **Cohere.**
Title: Simple Contrastive Learning of Sentence Embeddings.
- 2021/04 **ASAPP.**
Title: Simple Contrastive Learning of Sentence Embeddings.

Mentoring

- Anirudh Ajith (2024), MSE student at Princeton University. Now a predoctoral young investigator at AI2.
- Zhiyuan Zeng (2023), undergraduate student at Tsinghua University. Now a PhD student at UW.
- Jiatong Yu (2022-2023), undergraduate at Princeton University. Now an MSE student at Princeton.
- Howard Yen (2022-2024), undergraduate student at Princeton University. Now a PhD student at Princeton.
- Jane Pan (2021-2023), MSE student at Princeton University. Now a PhD student at NYU.
- Huihan Li (2020-2021), MSE student at Princeton University. Now a PhD student at USC.
- Xingcheng Yao (2021), undergraduate student at Tsinghua University. Now a PhD student at UCLA.
- Zichun Yu (2020), undergraduate student at Tsinghua University. Now a PhD student at CMU.
- Hao Peng (2019-2020), undergraduate student at Tsinghua University. Now a PhD student at Tsinghua.

Professional Services

Workshop Co-organizer

- [The First Workshop on Long Context Foundation Models \(LCFM\)](#). ICML 2024.

Area Chair

- Association for Computational Linguistics (ACL) Rolling Review: 2024.

Reviewer / Program Committee

- Association for Computational Linguistics (ACL) Rolling Review: 2022, 2023, 2024.
- Conference on Language Modeling (COLM): 2024.
- International Conference on Learning Representations (ICLR): 2023.
- Annual Meeting of the Association for Computational Linguistics (ACL): 2021.
- Conference on Empirical Methods in Natural Language Processing (EMNLP): 2021, 2022.
- Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL): 2021.
- Conference on Automated Knowledge Base Construction (AKBC): 2021.
- ICLR 2024 Workshop on Reliable and Responsible Foundation Models (R2-FM).
- ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models (ME-FoMo).

Departmental Services

- Compute management, Princeton Language and Intelligence (2023-2024).
- Compute management, Princeton NLP group (2020-2022).