

Tianyu Gao

tianyug@princeton.edu | gaotianyu.xyz/about

EDUCATION

Princeton University

Princeton, New Jersey, USA

Ph.D. Student in Computer Science

2020 - Present

- Research area: Natural Language Processing, Machine Learning
- Advisor: Prof. Danqi Chen

Tsinghua University

Beijing, China

B.Eng. in Computer Science and Technology

2016 - 2020

- Advisor: Prof. Zhiyuan Liu

FELLOWSHIPS, AWARDS AND HONORS

IBM PhD Fellowship, 2023.

The School of Engineering and Applied Science (SEAS) Award for Excellence, Princeton University, 2023.

Qualcomm Innovation Fellowship Finalists, 2023.

Outstanding Paper Award, ACL 2022.

Special-Class Scholarship (10 out of 3,000), Tsinghua University, 2019.

Cai Xiong Scholarship (10 out of 3,000), Tsinghua University, 2019.

Fellowship of the Initiative Scientific Research Program (10 out of 3,000), Tsinghua University, 2019.

PUBLICATIONS

- [1] *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.
- [2] *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.**
- [3] *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.
- [4] *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.
- [5] *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.
- [6] *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.**
- [7] *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.
- [8] *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.
- [9] *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.

- [10] *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.
- [11] *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.
- [12] *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.
- [13] *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.
- [14] *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.
- [15] *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.
- [16] *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.
- [17] *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.
- [18] *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.
- [19] *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.
- [20] *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.

PREPRINTS

- [1] *Sheared LLaMA: Accelerating Language Model Pre-training via Structured Pruning.*
Mengzhou Xia, **Tianyu Gao**, Zhiyuan Zeng, Danqi Chen. 2023.
- [2] *Evaluating Large Language Models at Evaluating Instruction Following.*
Zhiyuan Zeng, Jiatong Yu, **Tianyu Gao**, Yu Meng, Tanya Goyal, Danqi Chen. 2023.

* indicates equal contribution.

RESEARCH AND INDUSTRY EXPERIENCE

Facebook AI Research (FAIR) Research Intern with Jason Weston	Jun. 2022 – Aug. 2022 New York City, NY, USA
Tsinghua NLP Lab Research Assistant with Prof. Zhiyuan Liu	Nov. 2017 – May 2020 Beijing, China
Mila-Quebec AI Institute Research Assistant with Prof. Jian Tang	Jul. 2019 – Sept. 2019 Montreal, Canada

WeChat AI, Tencent*Tencent Rhino-Bird Elite Training Program with Peng Li*

May 2019 - May 2020

*Beijing, China***Momenta***Research Intern with Ji Liang*

May 2017 - May 2018

Beijing, China

TEACHING EXPERIENCE**COS484: Natural Language Processing***Teaching assistant. Instructor: Prof. Karthik Narasimhan*

Spring 2022

*Princeton University***COS324: Introduction to Machine Learning***Teaching assistant. Instructors: Prof. Sanjeev Arora and Prof. Danqi Chen*

Fall 2021

*Princeton University***Natural Language Processing***Teaching assistant. Instructor: Prof. Zhiyuan Liu*

Spring 2020

Tsinghua University

INVITED TALKS

- Semantic Machines, 2023. Enabling Large Language Models to Generate Text with Citations.
- Apple, 2023. Enabling Large Language Models to Generate Text with Citations.
- SingleStore, 2023. Enabling Large Language Models to Generate Text with Citations.
- MLNLP Seminar, 2023. Enabling Large Language Models to Generate Text with Citations (in Chinese).
- AI Time, 2022. Ditch the Gold Standard: Re-evaluating Conversational Question Answering (in Chinese).
- Cohere, 2021. Simple Contrastive Learning of Sentence Embeddings.
- ASAPP, 2021. Simple Contrastive Learning of Sentence Embeddings.

MENTORING

- Zhiyuan Zeng (2023), undergraduate student at Tsinghua University.
- Jiatong Yu (2022-2023), undergraduate student at Princeton University.
- Howard Yen (2022-2023), undergraduate student at Princeton University.
- 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.
- Zichun Yu (2020), undergraduate student at Tsinghua University.
- Hao Peng (2019-2020), undergraduate student at Tsinghua University. Now a PhD student at Tsinghua.

PROFESSIONAL SERVICE**Program Committee Member**

- Association for Computational Linguistics (ACL) Rolling Review.
- International Conference on Learning Representations (ICLR).
- Annual Meeting of the Association for Computational Linguistics (ACL).
- Conference on Empirical Methods in Natural Language Processing (EMNLP).
- Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL).
- International Conference on Computational Linguistics (COLING).
- Conference on Automated Knowledge Base Construction (AKBC).