

Yuxian Gu

Room 4-503, FIT Building, Tsinghua University, Beijing, China, 100084

☎ (+86) 15313594019 | ✉ guyx21@mails.tsinghua.edu.cn | 📠 t1101675

Education

Tsinghua University

PH.D. STUDENT IN COMPUTER SCIENCE AND TECHNOLOGY

- Member of **Tsinghua Conversational AI Group (CoAI)**
- Research Interest: Efficient Methods for Language Models
- Advisor: Prof. Minlie Huang

Beijing, China

Sept. 2021 - Present

Tsinghua University

B. ENG. IN COMPUTER SCIENCE AND TECHNOLOGY

Beijing, China

Sept. 2017 - June 2021

Selected Publications

- [1] **Jet-Nemotron: Efficient Language Model with Post Neural Architecture Search Architecture**
Yuxian Gu, Qinghao Hu, Shang Yang, Haocheng Xi, Junyu Chen, Song Han, Hai Cai. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2025.
- [2] **Data Selection via Optimal Learning for Language Models**
Yuxian Gu, Li Dong, Hongning Wang, Yaru Hao, Qingxiu Dong, Minlie Huang, Furu Wei. In *International Conference on Learning Representations (ICLR)*, 2025. **(Oral, top 1.8%)**
- [3] **Instruction Pre-Training: Language Models are Supervised Multitask Learners**
Daixuan Cheng, Yuxian Gu, Shaohan Huang, Junyu Bi, Minlie Huang, Furu Wei. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*.
- [4] **MiniLLM: Knowledge Distillation of Large Language Models**
Yuxian Gu, Li Dong, Furu Wei, Minlie Huang. *International Conference on Learning Representations (ICLR)*, 2024. **(560+ citations)**
- [5] **PPT: Pre-Trained Prompt Tuning for Few-Shot Learning**
Yuxian Gu*, Xu Han*, Zhiyuan Liu, Minlie Huang. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022. **(580+ citations)**
- [6] **Pre-Training to Learn in Context**
Yuxian Gu, Li Dong, Furu Wei, Minlie Huang. In *Proceedings of the 61th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023. **(Oral)**
- [7] **MiniPLM: Knowledge Distillation for Pre-training Language Models**
Yuxian Gu, Hao Zhou, Fandong Meng, Jie Zhou, Minlie Huang. *International Conference on Learning Representations (ICLR)*, 2025.
- [8] **Learning Instructions with Unlabeled Data for Zero-Shot Cross-Task Generalization**
Yuxian Gu, Pei Ke, Xiaoyan Zhu, Minlie Huang. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022. **(Oral)**
- [9] **CPM-2: Large-Scale Cost-Effective Pre-Trained Language Models**
Zhengyan Zhang*, Yuxian Gu*, Xu Han*, Shengqi Chen*, Chaojun Xiao*, Zhenbo Sun, Yuan Yao, et al. In *AI Open*, 2022.
- [10] **Train No Evil: Selective Masking for Task-Guided Pre-Training**
Yuxian Gu, Zhengyan Zhang, Xiaozhi Wang, Zhiyuan Liu, Maosong Sun. In *Proceedings of the Empirical Methods in Natural Language Processing (EMNLP)*, 2020. (short)

Experience

MIT HAN Lab

VISITING STUDENT

- Advised by Prof. Song Han.

Cambridge, USA

Oct. 2024 - Jun. 2025

Microsoft Research Asia

RESEARCH INTERN

- Mentored by Dr. Li Dong.

Beijing, China

Oct. 2022 - Mar. 2024

Honors & Awards

2025	National Scholarship Awarded to Chinese Ph.D. students with excellent achievements.	<i>Tsinghua University</i>
2025	In-Tech Scholarship from Ant Research Encourages Chinese doctoral students worldwide to solve the most valuable technical problems in computer science. (Awarded to the top 10 Chinese students)	<i>Ant Research</i>
2025	Apple Scholars in AIML PhD fellowship Recognizes the contributions of researchers in computer science and engineering at the graduate and postgraduate levels. (Awarded to the top 21 students worldwide)	<i>Apple Inc.</i>
2022-2024	Comprehensive Scholarship Awarded to Ph.D. students with excellent achievements by Tsinghua University.	<i>Tsinghua University</i>
2021	Excellent Graduate, Tsinghua University (top 5%) Awarded to graduates with excellent achievements by Tsinghua University.	<i>Tsinghua University</i>
2021	Outstanding Graduate, Dept. CST, Tsinghua University Awarded to graduates with excellent achievements by Department of Computer Science and Technology, Tsinghua University.	<i>Tsinghua University</i>
2021	Outstanding Undergraduate Dissertation, Tsinghua University (top 5%) Awarded to outstanding undergraduate dissertations by Tsinghua University.	<i>Tsinghua University</i>
2019	Silver Prize, 2019 Asia Student Supercomputer Challenge World-wide contest on super-computing.	<i>Dalian, China</i>
2018, 2020	Comprehensive Scholarship Awarded to students with excellent achievements by Department of Computer Science and Technology, Tsinghua University.	<i>Tsinghua University</i>
2016	1st Prize, Chinese Physics Olympiad in Provinces Nation-wide contest on physics.	<i>Beijing, China</i>
2016	1st Prize, Pan-Pearl River Delta and Chinese Elite Schools Physics Olympiad Nation-wide contest on physics.	<i>Shenzhen, China</i>

Projects

Jet-Nemotron: SOTA Hybrid Model with Linear and Full Attention

DEC. 2024 - JULY. 2025

- Matches Qwen3 on MMLU(-pro), math, code, and long-context tasks. Delivering **>40x** generation throughput.
- Open Source: <https://github.com/NVlabs/Jet-Nemotron>

CPM-2: Large-scale Cost-effective Pre-trained Language Models

JAN. 2021 - JUNE. 2021

- (By the time of release) The largest open-source Chinese pre-trained language model (**11B parameters**).
- Open Source: <https://github.com/TsinghuaAI/CPM-2-Pretrain>

EVA: Large-scale Open-domain Chinese Dialogue System

SEP. 2021 - JAN. 2022

- (By the time of release) The largest open-source Chinese dialogue model with pre-training (2.8B parameters).
- Open Source: <https://github.com/thu-coai/EVA>

Teaching Assistantships

Object Oriented Programming

Instructor: Prof. Minlie Huang

Tsinghua University
Spring 2022, 2023, 2024

Artificial Neural Networks

Instructor: Prof. Minlie Huang

Tsinghua University
Fall 2021, 2022, 2023