

Yanggan Gu

yanggan@outlook.com | [Google Scholar](#) | [Homepage](#) | [GitHub](#)

Education

The Hong Kong Polytechnic University , Hong Kong, China. <i>Ph.D. student in Computing.</i>	2026/01 - Present Advisor: Prof. Hongxia Yang
Soochow University [<i>Project 211</i>], Suzhou, China. <i>M.S. in Computer Science and Technology.</i>	2022/09 - 2025/06 GPA 3.7 (Top-5%)
Guangdong University of Technology , Guangzhou, China. <i>B.S. in Software Engineering.</i>	2018/09 - 2022/06 GPA 4.0 (Top-1%)

Research Experience

Model Fusion in LLMs , <i>Researcher</i> Department of Computing, Hong Kong Polytechnic University . <i>Research Topic:</i> Model Fusion, Alignment of LLMs, Reasoning in LLMs	2025/04 - Present Advisor: Prof. Hongxia Yang
Trustworthy LLMs , <i>Research Assistant</i> AI Thrust, The Hong Kong University of Science and Technology (Guangzhou). <i>Research Topic:</i> Alignment of LLMs, Factual Reasoning in LLMs, Knowledge Distillation	2024/06 - 2025/03 Advisor: Prof. Xuming Hu
Syntactic Parsing , <i>Master's Researcher</i> School of Computer Science and Technology, Soochow University . <i>Research Topic:</i> Linguistic Structure Prediction, Text2SQL (maintaining SeSQL dataset)	2022/09 - 2025/06 Advisor: Prof. Zhenghua Li

Selected Publications

Model Merging Scaling Laws in Large Language Models [Code] <i>Co-first Author</i>	ICML 2026 (CCF-A)
Develops scaling laws that predict model-merging performance across model sizes, expert counts, and data regimes.	
From Parameters to Behaviors: A Survey of Model Fusion for Large Language Models <i>Co-first Author</i>	EMNLP 2026 Findings (CCF-B)
Provides a unified taxonomy of model fusion spanning parameter-, prediction-, and behavior-level methods.	
InfiFPO: Implicit Model Fusion via Preference Optimization in Large Language Models [Code] <i>First Author</i>	NeurIPS 2025 Spotlight (CCF-A)
Introduces sequence-level probability fusion through preference optimization and validates it on 11 benchmarks.	
Capturing Nuanced Preferences: Preference-Aligned Distillation for Small Language Models [Code] <i>First Author</i>	ACL 2025 Findings (CCF-A)
Distills full preference distributions into small language models for more precise preference alignment.	
Reasoning Factual Knowledge in Structured Data with Large Language Models [Code] <i>Co-first Author</i>	ACL 2025 Findings (CCF-A)
Introduces StructFact, a 13,407-query benchmark for reasoning over heterogeneous structured knowledge.	

Honors and Awards

• 1st place in the CCL-2023 Chinese AMR Parsing Evaluation Task. [Report]	2023/08
• Top 10 Innovative Figures of Guangdong University of Technology. [Sohu News]	2021/12
• National Encouragement Scholarship. [Link]	2021/11

Academic Service

- **Reviewer:** ICML, ICLR, NeurIPS, ARR, and TMLR.