



Yifan Zhang

☎ (+1) 310-339-2010 ✉ yif-zhang@outlook.com  [yifzhang.com](https://www.linkedin.com/in/yifzhang.com)  github.com/yifanzhang-pro

Education

Department of Computer Science, University of California, Los Angeles <i>Visiting Graduate Student in Computer Science</i> <ul style="list-style-type: none">• Advisor: Quanquan Gu	2024 - 2025 Los Angeles, CA
Institute for Interdisciplinary Information Sciences (IIIS), Tsinghua University <i>Graduate Student (MPhil) in Computer Science</i> <ul style="list-style-type: none">• Advisors: Andrew Chi-Chih Yao, Yang Yuan	2021 - 2024 Beijing
Yuanpei College, Peking University <i>B.S. in Data Science & Computer Science (Top 2 / 36)</i> <ul style="list-style-type: none">• Advisors: Zihua Zhang, Zhouchen Lin	2017 - 2021 Beijing
Experimental High School Attached to Beijing Normal University (ESBNU) <i>High School Diploma</i> <ul style="list-style-type: none">• High School Student	2011 - 2017 Beijing

Selected Publications (Google Scholar)

- Augmenting Math Word Problems via Iterative Question Composing (AAAI 2025)**
 - Haoxiong Liu*, Yifan Zhang*, Yifan Luo, Andrew Chi-Chih Yao
- General Preference Modeling with Preference Representations for Aligning Language Models (ICLR 2025 submission)**
 - Yifan Zhang*, Ge Zhang*, Yue Wu*, Kangping Xu, Quanquan Gu
- Beyond Squared Error: Exploring Loss Design for Enhanced Training of Generative Flow Networks (ICLR 25 submission)**
 - Rui Hu*, Yifan Zhang*, Zhuoran Li, Longbo Huang (Rating 886 on Openreview)
- Cumulative Reasoning for Large Language Models (BGPT @ ICLR 2024)**
 - Yifan Zhang*, Jingqin Yang*, Yang Yuan, Andrew Chi-Chih Yao
- Information Flow in Self-Supervised Learning (ICML 2024)**
 - Zhiqian Tan, Jingqin Yang, Weiran Huang, Yang Yuan†, Yifan Zhang†
- Matrix Information Theory for Self-Supervised Learning (ICML 2024)**
 - Yifan Zhang*, Zhiqian Tan*, Jingqin Yang*, Weiran Huang, Yang Yuan
- Contrastive Learning Is Spectral Clustering on Similarity Graph (ICLR 2024)**
 - Zhiqian Tan*, Yifan Zhang*, Jingqin Yang*, Yang Yuan
- Trade-off Between Efficiency and Consistency for Removal-based Explanations (NeurIPS 2023)**
 - Yifan Zhang*, Haowei He*, Zhiqian Tan, Yang Yuan

Publications & Manuscripts (Full List)

- Augmenting Math Word Problems via Iterative Question Composing (AAAI 2025)**
 - Haoxiong Liu*, Yifan Zhang*, Yifan Luo, Andrew Chi-Chih Yao
- General Preference Modeling with Preference Representations for Aligning Language Models (ICLR 2025 submission)**
 - Yifan Zhang*, Ge Zhang*, Yue Wu*, Kangping Xu, Quanquan Gu
- Beyond Squared Error: Exploring Loss Design for Enhanced Training of Generative Flow Networks (ICLR 25 submission)**
 - Rui Hu*, Yifan Zhang*, Zhuoran Li, Longbo Huang (Rating 886 on Openreview)
- Scaling Image Tokenizers with Grouped Spherical Quantization (CVPR 2025 submission)**
 - Jiangtao Wang, Zhen Qin, Yifan Zhang, Tao Hu, Björn Ommer, Rania Briq, Stefan Kesselheim
- Matrix Information Theory for Self-Supervised Learning (ICML 2024)**
 - Yifan Zhang*, Zhiqian Tan*, Jingqin Yang*, Weiran Huang, Yang Yuan

Information Flow in Self-Supervised Learning (ICML 2024)

- [Zhiquan Tan](#), [Jingqin Yang](#), [Weiran Huang](#), [Yang Yuan](#)[†], [Yifan Zhang](#)[†]

Cumulative Reasoning for Large Language Models (BGPT @ ICLR 2024)

- [Yifan Zhang](#)^{*}, [Jingqin Yang](#)^{*}, [Yang Yuan](#), [Andrew Chi-Chih Yao](#)

Meta Prompting for AI Systems (BGPT @ ICLR 2024)

- [Yifan Zhang](#), [Yang Yuan](#), [Andrew Chi-Chih Yao](#)

AutoMathText: Autonomous Data Selection with Language Models for Mathematical Texts (DPFM @ ICLR 2024)

- [Yifan Zhang](#)^{*}, [Yifan Luo](#)^{*}, [Yang Yuan](#), [Andrew Chi-Chih Yao](#)

Contrastive Learning Is Spectral Clustering on Similarity Graph (ICLR 2024)

- [Zhiquan Tan](#)^{*}, [Yifan Zhang](#)^{*}, [Jingqin Yang](#)^{*}, [Yang Yuan](#)

SEAL: Simultaneous Label Hierarchy Exploration and Learning (TMLR 2024)

- [Zhiquan Tan](#)^{*}, [Zihao Wang](#)^{*}, [Yifan Zhang](#)^{*}

EffCause: Discover Dynamic Causal Relationships Efficiently from Time-Series (TKDD 2024)

- [Yicheng Pan](#), [Yifan Zhang](#), [Xinrui Jiang](#), [Meng Ma](#), [Ping Wang](#)

Coded real number matrix multiplication for on-device edge computing (SPL 2023)

- [Zhiquan Tan](#), [Dingli Yuan](#), [Yifan Zhang](#), [Zhongyi Huang](#)

Trade-off Between Efficiency and Consistency for Removal-based Explanations (NeurIPS 2023)

- [Yifan Zhang](#)^{*}, [Haowei He](#)^{*}, [Zhiquan Tan](#), [Yang Yuan](#)

Preprints

Training and Evaluating Language Models with Template-based Data Generation

- [Yifan Zhang](#), *et al.*

FlashSampling: Fast and Memory-Efficient Exact Sampling with Group-Gumbel-Max

- [Zhen Qin](#), [Xuyang Shen](#), [Yifan Zhang](#), [Yiran Zhong](#)

On the Diagram of Thought

- [Yifan Zhang](#), [Yang Yuan](#), [Andrew Chi-Chih Yao](#)

Experience

Research on LLM Architecture Design & LLM Alignment Algorithms

University of California, Los Angeles

09/2024 - now

- We developed a novel approach, the General Preference representation Model (GPM), that embeds human preferences into a latent space for efficient and expressive preference learning. We introduced General Preference Optimization (GPO) to extend reward-based reinforcement learning from human feedback. Demonstrated significant performance improvements over traditional models in preference modeling, cyclic preferences, and downstream tasks (as Huggingface daily papers) (ICLR 2025 submission).
- (Ongoing) We propose sequence modeling architectures with linear complexities with theoretical guarantees.
- (Ongoing) Investigating how to effectively perform token-level LLM alignment, especially address the chosen log-likelihood decreasing phenomenon in DPO and its variants.

Research on AI Reasoning & Reinforcement Learning

IIS, Tsinghua University & University of California, Los Angeles

03/2023 - Now

- We propose the *Cumulative Reasoning* (CR) framework for large language models (LLMs), a novel approach to complex reasoning. This framework introduces multiple agents functioning as proposers, verifiers, and reporters, achieving state-of-the-art performance on several reasoning benchmarks, including the Game of 24, MATH, and logical inference tasks (BGPT @ ICLR 2024).
- Proposed *Meta Prompting*, an innovative technique reshaping the utilization of large language models (LLMs) and multi-modal foundation models in problem-solving and data interaction. Inspired by type theory, meta prompting emphasizes the structure and syntax of examples over their detailed content (BGPT @ ICLR 2024).

- Developed the *Iterative Question Composing (IQC)* method for mathematical problem generation, enabling infinite data synthesis. Empirically, open-source models fine-tuned with data generated by this method achieve SOTA performance (AAAI 2025).
- Introduced *Autonomous Data Selection (AutoDS)* with language models as generative verifiers (GenRM) for data selection and released the AutoMathText dataset on Huggingface, which once trended first on the leaderboard (as Huggingface daily papers) (DPFM @ ICLR 2024).
- We introduced the Diagram of Thought (DoT), a framework that models reasoning as constructing a directed acyclic graph (DAG) within a single large language model. This approach allows for iterative refinement of propositions, critiques, and verifications, improving logical consistency and reasoning. Formalized using the Topos theory, DoT enables enhanced, structured reasoning without needing external control mechanisms.
- Proposed new regression losses to improve Generative Flow Networks (GFlowNets), focusing on better sample diversity, faster convergence, and increased robustness. Demonstrated superior performance in tasks such as molecule generation (ICLR 2025 submission).

Research on Representation Learning & Interpretability

IIIS, Tsinghua University

09/2021 - 09/2023

- Investigated the theoretical framework of interpretability and proved an impossible trinity theorem for attribution methods (NeurIPS 2023).
- Explored the underlying mechanisms behind contrastive learning, demonstrating the exact equivalence between standard contrastive learning and spectral clustering (ICLR 2024).
- Established a matrix-based theoretical framework for analyzing the representation of pre-trained models, incorporating matrix-based learning objectives leading to downstream performance improvements on several tasks (ICML 2024).

Skills

English Proficiency: IELTS 7.5, GRE 329

Programming Languages: Python, C++, C, Matlab, Bash, Scala

Technologies: Pytorch, Linux, Docker, Kubernetes

Maths: Familiar with mathematical analysis, functional analysis, differential geometry, linear algebra, category theory

Awards: 2nd Prize, National High School Physics Olympics (CPhO); May 4th Scholarship at Peking University

Professional Activities

Conference Reviewer for ICLR, ICML, NeurIPS, AAAI, AISTATS

Journal Reviewer for TKDD, INS, NEUNET, NEUCOM, JER, JOLT, COR, MAVIC

Invited Talks

On General Preference Modeling with Preference Representations (slides), November 2024, **Bytedance AI Seminar**

On Cumulative Reasoning with Large Language Models (slides), September 2023, **BAAI**

On Contrastive Learning and Spectral Clustering (slides), May 2023, **Microsoft Research, Asia**