Yuzhen Huang

Email: yhuanghj@cse.ust.hk Homepage: hyz17.github.io

Research Interests: I am primarily focused on large language models, particularly in advancing their reasoning capabilities and multimodal understanding. To achieve this, my research interests lie in: (1) enhancing reasoning, planning, and agentic abilities through self-improvement and RL techniques (2) developing reliable evaluation methods for language models.

Education

The Hong Kong University of Science and Technology, Hong Kong SAR, China. *Feb. 2024 - present*
– **Ph.D in Computer Science**
– Advisor: Prof. Junxian He

Shanghai Jiao Tong University, Shanghai, China *Sep. 2019 – Jul. 2023*
– **B.Eng. in Computer Science**
– *GPA*: 3.89/4.3, *Score*: 90.27/100

Research Projects

[1] SimpleRL-Zoo: Investigating and Taming Zero Reinforcement Learning for Open Base Models in the Wild
– **Achieve improvements in both reasoning accuracy and response length across diverse models.**
– Introduce SimpleRL-Zoo, a simple reinforcement learning recipe to improve models' reasoning abilities.
– Identify key factors that shape the emergence of advanced reasoning behaviors (i.e., the "aha moment").
– Published in COLM 2025 as the co-first author.

[2] B-STaR: Monitoring and Balancing Exploration and Exploitation in Self-Taught Reasoner
– **Quantitatively analyze the dynamics of exploration and exploitation during self-improvement.**
– Introduce B-STaR, a Self-Taught Reasoning framework that autonomously adjusts its configurations.
– Balance exploration and exploitation, leading to superior performance.
– Published in ICLR 2025 as the co-first author.

[3] Compression Represents Intelligence Linearly
– **Investigate the linear correlation between compression and intelligence in LLMs.**
– Provide evidence for the belief that superior compression is indicative of greater intelligence.
– Propose compression efficiency serves as an unsupervised and reliable metric to assess LLMs' abilities.
– Published in COLM 2024 as the first author.

[4] C-Eval: A Multi-Level Multi-Discipline Chinese Evaluation Suite for Foundation Models
– **The first comprehensive Chinese evaluation suite for LLMs.**
– Conduct a thorough evaluation of the most advanced LLMs.
– Over 9.8M downloads on Hugging Face and more than 100 models on the leaderboard.
– Published in NeurIPS 2023 as the first author.

Publications

* **denotes equal contribution**

[1] J Li*, W Zhao*, J Zhao*, W Zeng*, H Wu*, X Wang, R Ge, Y Cao, **Y Huang**, W Liu, J Liu, Z Su, Y Guo, F Zhou, L Zhang, J Michelini, X Wang, X Yue, S Zhou, G Neubig, J He. The Tool Decathlon: Benchmarking Language Agents for Diverse, Realistic, and Long-Horizon Task Execution. Arxiv.

[2] **Y Huang***, W Zeng*, X Zeng, Q Zhu, J He. From Accuracy to Robustness: A Study of Rule- and Model-based Verifiers in Mathematical Reasoning. Arxiv.

[3] W Liu, R Zhou, Y Deng, **Y Huang**, J Liu, Y Deng, Y Zhang, J He. Learn to Reason Efficiently with Adaptive Length-based Reward Shaping. Arxiv.

[4] W Zeng*, **Y Huang***, Q Liu*, W Liu, K He, Z Ma, J He. SimpleRL-Zoo: Investigating and Taming Zero Reinforcement Learning for Open Base Models in the Wild. COLM 2025

[5] K Shum*, **Y Huang***, H Zou, D Qi, Y Liao, X Chen, Q Liu, J He. Predictive Data Selection: The Data That Predicts Is the Data That Teaches. ICML 2025.

[6] W Zang*, **Y Huang***, L zhao, Y Wang, Z Shan, J He. B-STaR: Monitoring and Balancing Exploration and Exploitation in Self-Taught Reasoner. ICLR, 2025.

[7] **Y Huang***, J Zhang*, Z Shan, J He. Compression Represents Intelligence Linearly. Conference on Language Modeling (COLM), 2024.

[8] **Y Huang***, Y Bai*, Z Zhu, J Zhang, J Zhang, T Su, J Liu, C Lv, Y Zhang, Y Fu, M Sun, J He. C-Eval: A Multi-Level Multi-Discipline Chinese Evaluation Suite for Foundation Models. NeurIPS (Datasets and Benchmarks track), 2023

Internship

Research Intern, Tencent *Nov. 2023 – Jan. 2024*
Mentor: Zifei Shan

Intern Research Scientist, TikTok Singapore *Jan. 2025 – Now*
Mentor: Dr. Qian Liu

Professional Activities

Reviewer: NeurIPS 2024, ICLR 2025, ICML 2025, NeurIPS 2025 (Top Reviewer), ICLR 2026, ARR

Talks

[1] Emerging Reasoning with Reinforcement Learning is Both Effective and Efficient. *Mar 2025, Georgia Tech PAIR*

[2] Emerging Reasoning with Reinforcement Learning is Both Effective and Efficient. *Feb 2025, Apple AIML*

[3] Compression Represents Intelligence Linearly. *May 2024, BAAI*

Teaching

Teaching Assistant, The Hong Kong University of Science and Technology *Fall 2024*
COMP5212 Machine Learning

Teaching Assistant, The Hong Kong University of Science and Technology *Fall 2025*
COMP4901B Large Language Models

Standard Tests

[1] **TOEFL** – 102