

Yi Hu

huyi2002@stu.pku.edu.cn | [Google Scholar](#) | [GitHub](#) | [Personal Website](#)

Education

Peking University	Sep 2024 - Jul 2029 (Expected)
PhD in Artificial Intelligence	
Coursework: Machine Learning, AI Reasoning, Large Language Model Alignment, Computer Vision, AI Arts	
Peking University	Sep 2020 - Jul 2024
BSc in Physics	
Coursework: Advanced Mathematics, Linear Algebra, Data Structure and Algorithm, Methods of Mathematical Physics, Theoretical Mechanics, Quantum Mechanics	
Thesis: Classification of Flow and Nonflow Using Machine Learning	

Research Interests

Personal Website

My research seeks to define the boundaries of large language models reasoning, revealing what can be reliably expected from them and what remains beyond their reach. I systematically investigate how LLMs reason, *where they fail, why failures occur*, and *how to fix them*. Specifically, my work focuses on:

- *Diagnosing systematic failure points* in LLM reasoning across mathematical, numerical, and physical domains;
- *Uncovering underlying mechanisms*, such as case-based / rule-based reasoning patterns and depth under-utilization;
- *Designing targeted improvements*, including training strategies that enhance strict rule-following and strengthen length generalization across varying tasks.

Publications

Google Scholar

First-/Co-first-author papers:

[1] **Beyond Single-Task: Robust Multi-Task Length Generalization for LLMs** *NeurIPS 2025*
Yi Hu*, Shijia Kang*, Haotong Yang, Haotian Xu, Muhan Zhang†.

[2] **What Affects the Effective Depth of Large Language Models?** *NeurIPS 2025 MI Workshop*
Yi Hu, Cai Zhou, Muhan Zhang†.

[3] **Case-Based or Rule-Based: How Do Transformers Do the Math?** *ICML 2024*
Yi Hu, Xiaojuan Tang, Haotong Yang, Muhan Zhang†.

[4] **Code prompting: a neural symbolic method for complex reasoning in large language models** *preprint 2023*
Yi Hu, Haotong Yang, Zhouchen Lin, Muhan Zhang†

Selected Co-authored papers:

[1] **Coevolutionary continuous discrete diffusion: Make your diffusion language model a latent reasoner?** *submitted to ICLR 2026*
Cai Zhou, Chenxiao Yang, Yi Hu, Chenyu Wang, Chubin Zhang, Muhan Zhang, Lester Mackey, Tommi Jaakkola, Stephen Bates, Dinghuai Zhang†

[2] **Phybench: Holistic evaluation of physical perception and reasoning in large language models?** *NeurIPS 2025*
Shi Qiu*, Shaoyang Guo*, Zhuo-Yang Song*, Yunbo Sun*, Zeyu Cai*, Jiashen Wei*, Tianyu Luo*, Yixuan Yin, Haoxu Zhang, Yi Hu, and others

[3] **Number Cookbook: Number Understanding of Language Models and How to Improve It** *ICLR 2025*
Haotong Yang, Yi Hu, Shijia Kang, Zhouchen Lin, Muhan Zhang†

Experience

LLM Post-Training Team, Xiaohongshu (RedNote / Little Red Book)

Beijing, China

Research Intern

Oct 2024 – Jun 2025

- **Research Focus:** Following o1; investigating how to enhance LLM reasoning and reproduce test-time scaling laws.
- **Responsibilities:** Collected and evaluated benchmarks for logical reasoning; assessed reasoning models; constructed long-chain reasoning datasets and built an RL environment. Experimented with multiple post-training techniques, including SFT, DPO, and PPO, to improve model performance on logical reasoning tasks.
- **Achievements:** Developed a specialized dataset for logical reasoning and enhanced the base model (Qwen-32B), raising its logical reasoning performance to a level comparable with QwQ. Contributed to the [technical report](#).

LLM Research Group, Huawei

Beijing, China

University-Industry Collaboration Project

Dec 2024 – Dec 2025

- **Research Focus:** Contributed to [Align-Anything](#) (4.6k stars), a multimodal post-training framework, ensuring compatibility with Huawei's Ascend AI processors.
- **Responsibilities:** Integrated multi-turn RL functions from [RAGEN](#) into the Align-Anything framework; refactored core components including environment interaction, reward computation, and policy update logic, and successfully integrated them with the framework's RL trainer; ensured its compatibility with the Huawei Ascend 910B NPU.
- **Achievements:** Successfully implemented the multi-turn RL module within Align-Anything, replicating performance consistent with the original methodology.

School of Physics, Peking University

Beijing, China

Research Intern (Supervisor: Prof. [Huichao Song](#))

Nov 2021 – Nov 2022

- **Research Focus:** Classifying flow and nonflow signals in Quark Gluon Plasma (QGP), with potential to enhance experimental data purity in high-energy physics.
- **Responsibilities:** Collaborated in designing and implementing XGBoost as the classification model; analyzed and preprocessed QGP simulation datasets; optimized hyperparameters to improve classification performance.
- **Achievements:** Delivered a well-performing classification model that effectively distinguishes between flow and nonflow signals, demonstrating strong potential for future application in experimental data purification. This research was subsequently developed into my undergraduate thesis.

Selected Honors and Awards

Outstanding Graduate Award, Peking University

Jun 2024

Ubiquant Scholarship, Peking University

Sep 2023

Xiaomi Scholarship, Peking University

Sep 2022

Technical & Research Skills

- **Model Expertise:** Familiar with open-source LLMs (e.g., LLaMA, Qwen); skilled in post-training techniques including SFT, DPO, PPO and GRPO; familiar with mechanistic interpretability methods, such as activation steering and linear probing.
- **Programming:** Proficient in Python; skilled in PyTorch and Huggingface Transformers. Contributed to open-source projects like [Align-Anything](#) (4.6k stars).
- **Research:** Led multiple projects resulting in publications at top-tier conferences (ICML, NeurIPS). Experienced in the full research lifecycle: idea formulation, experimental design, paper writing, and peer-review process.