

## Education

- 2023–2027 **Renmin University of China** – B.Eng. Computer Science (Turing Elite Program) + B.Econ. Finance  
Beijing, China **GPA 3.94/4.00 · Rank 1/30 · Research Area Large Language Models**  
**Selected Coursework** Calculus (96), Adv. Linear Algebra (96), Probability (94), Computer Systems (95)
- 2026.01 **University of California, Berkeley** – Visiting Student · Berkeley NLP / BAIR  
– 2026.08 **Advanced NLP (CS 288)** – **A+**, *highest score in class (1/150)* · Taught by *Sewon Min, Alane Suhr*  
Berkeley, CA **Computational Models of Language Use (CS 294-258)** – **A+** (top 4/25) · Taught by *Alane Suhr*

## Publications

- Preprint* **Which Models Are Our Models Built On? Auditing Invisible Dependencies in Modern LLMs**  
2026 Sanjay Adhikesaven\*, Haoxiang Sun\*, Sewon Min\* <sup>equal contribution</sup> [arXiv · Code · Demo]  
We present *ModSleuth* – an agentic system that recursively reconstructs LLM dependency graphs from public artifacts (model & dataset cards, technical reports, code). Recovered **1,060** source-verified dependencies across four major releases, surfacing underreporting, hidden license paths, and train–eval coupling.
- ACL 2026 **Challenging the Boundaries of Reasoning: An Olympiad-Level Math Benchmark for LLMs**  
Oral Haoxiang Sun, Yingqian Min, Zhipeng Chen, Wayne Xin Zhao†, Ji-Rong Wen [arXiv · Code · Dataset]  
(67 cites) We present *OlymMATH* – the first bilingual parallel Olympiad-level benchmark: **200** verifiable problems (English/Chinese) plus **150** Lean formal proofs, sourced from print materials with per-problem expert review. Exposed guessing-based solving in SOTA models (DeepSeek-R1, o3-mini, Gemini 2.5 Pro).
- Preprint* **Enhancing LLM Reasoning with Reward-guided Tree Search**  
2024 Jinhao Jiang, ..., Haoxiang Sun, Jia Deng, Wayne Xin Zhao†, et al. [arXiv · Code]  
(69 cites) We present *STILL-1* – reward-guided tree search; built the Beam-Search baseline and a *sympy* math verifier.
- ACL 2024 **LLMBox: A Comprehensive Library for Large Language Models**  
Demo Tianyi Tang, Yiwen Hu, ..., Haoxiang Sun, ..., Wayne Xin Zhao†, Ji-Rong Wen [arXiv · Anthology · Code]  
(9 cites) We present *LLMBox* – a unified LLM training/eval platform; built evaluation and few-shot prompting.

## Research Experience

- 2026–Now **NLP Group / BAIR, UC Berkeley** – Research Intern · Advisor: [Sewon Min](#)  
Co-first-authored *ModSleuth*; built the agentic recursive graph-reconstruction pipeline.
- 2025–2026 **MLL Lab, Northwestern University** – Research Intern · Advisor: [Manling Li](#) · Mentor: [Zihan Wang](#)  
**RAGEN (2.7k\*)**: multi-turn agent context management + Lean env; RL training via *verL*.  
**Chain-of-Experts** (cyclic MoE): rewrote the MoE-variant pretraining & load-balancing pipeline.
- 2024–2025 **Gaoling School of AI, Renmin University of China** – Research Intern · Advisor: [Wayne Xin Zhao](#)  
Led *OlymMATH* (first author, ACL 2026 Oral); contributed to *LLMBox*.
- 2024 **Beijing Academy of Artificial Intelligence (BAAI)** – Research Intern · IR & Knowledge Computing  
LLM math reasoning; built the Beam-Search baseline and verifier for *STILL-1*.

## Honors & Awards (Selected)

- 2026 **SenseTime Scholarship** – 30 awardees nationwide
- 2024, 2025 **National Scholarship (x2)** – Ministry of Education, P.R. China (top ~0.2% nationwide)
- 2022 **National Silver Medal** – China Mathematical Olympiad (CMO); ranked 2nd in Anhui Province
- 2021 **National Gold Medal** – China Southeast Mathematical Olympiad (CSMO)
- 2024 **Asia Team Champion** – Asia Terminal, hosted by Citadel | Citadel Securities

## Activities (Selected)

- 2025.12 **RUC “Bairong Cup” AI Agent Contest** – **1st Place**; multimodal interview robot (Technical Lead)
- 2025.03 **ETH Hangzhou Hackathon** – MCP multi-agent on-chain data analysis & trading [Demo]
- 2025.03 **Discover Citadel | Citadel Securities** – visited HK office; Team Champion, Cryptography Challenge
- 2024–2025 **Beijing College Student Innovation Program** – Technical Lead, News-Image Generation  
Teacher–student pipeline: Qwen2.5-72B data synthesis, then LoRA + PPO on a 7B student model.