

SHIBO HAO

☎ 858-789-3662 ✉ s5hao@ucsd.edu 🌐 ber666.github.io 🏠 github.com/ber666 🎓 [Google Scholar](https://scholar.google.com/citations?user=ber666)

Education

University of California, San Diego <i>Ph.D. in Data Science. Advisor: Zhiting Hu</i>	Sep. 2022 – Present <i>La Jolla, CA</i>
Peking University <i>B.S. in Computer Science</i>	Sep. 2018 – Jun. 2022 <i>Beijing, China</i>

Experience

MBZUAI Institute of Foundation Models <i>Research Intern, post-training of a 70B fully open-source LLM</i>	Oct. 2025 – Present <i>Remote</i>
Bloomberg L.P. <i>Data Scientist Intern (Bloomberg Data Science Fellow), web agents</i>	Jun. 2025 – Sep. 2025 <i>New York, NY</i>
Meta FAIR Lab <i>Research Scientist Intern (Mentors: Yuandong Tian, Jason Weston), latent reasoning</i>	Jun. 2024 – Jan. 2025 <i>Menlo Park, CA</i>

Research Interests

My research focuses on pushing the boundaries of machine reasoning in large language models. I develop latent-space reasoning methods ([COCONUT](#), [COCONUT-Theory](#), [COCONUT-Dynamics](#)), train LLMs to reason via reinforcement learning ([Guru](#), [OREO](#), [Scaling](#)), build system-2 reasoning frameworks with world-model planning ([RAP](#), [LLM Reasoners](#), [Pandora](#)), and find better ways for LLM agents to interact with the world ([ToolkenGPT](#), [CocoaBench](#)). I also work on large-scale post-training of LLMs for general reasoning ability ([K2-V2](#), [K2-Think](#)).

Awards

Bloomberg Data Science Ph.D. Fellowship (3 recipients worldwide)	2024
Best Paper Award at SoCal NLP Symposium	2023
Outstanding Graduate Award, Peking University	2022
Leo KoGuan Scholarship (Top 1% at Peking University)	2022
First Prize, National Olympiad in Informatics in Provinces (NOIP)	2017

Selected Publications

(* indicates equal contribution)

Training Large Language Models to Reason in a Continuous Latent Space

Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, Yuandong Tian
COLM 2025

Highlighted by [Quanta Magazine](#)

Reasoning with Language Model is Planning with World Model

Shibo Hao*, Yi Gu*, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, Zhiting Hu
EMNLP 2023

Featured in [State of AI Report 2023](#)

ToolkenGPT: Augmenting Frozen Language Models with Massive Tools via Tool Embeddings

Shibo Hao, Tianyang Liu, Zhen Wang, Zhiting Hu

NeurIPS 2023 (Oral, 67 / 12345)

Best Paper Award at [SoCalNLP 2023](#)

LLM Reasoners: New Evaluation, Library, and Analysis of Step-by-Step Reasoning with Large Language Models

Shibo Hao*, Yi Gu*, Haotian Luo*, Tianyang Liu, Xiyan Shao, Xinyuan Wang, Shuhua Xie, Haodi Ma, Adithya Samavedhi, Qiyue Gao, Zhen Wang, Zhiting Hu

COLM 2024

2.3k Stars (as of Dec. 2025) at [Github](#)

Offline Reinforcement Learning for LLM Multi-Step Reasoning

Huaijie Wang*, **Shibo Hao***, Hanze Dong, Shenao Zhang, Yilin Bao, Ziran Yang, Yi Wu

Findings of ACL 2025

Also accepted to ICLR 2025 Workshop on Reasoning and Planning for LLMs (Oral, 7 / 181)

IsoCompute Playbook: Optimally Scaling Sampling Compute for RL Training of LLMs

Zhoujun Cheng*, Yutao Xie*, Yuxiao Qu*, Amrith Setlur*, **Shibo Hao**, Varad Pimpalkhute, Tongtong Liang, Feng Yao, Hector Liu, Eric Xing, Virginia Smith, Ruslan Salakhutdinov, Zhiting Hu, Taylor Killian, Aviral Kumar
Blog Post, <https://compute-optimal-rl-llm-scaling.github.io/>

CocoaBench: An Evaluation Framework for General Agents with Compositional Cognitive Abilities

Shibo Hao*, Zhining Zhang*, Zhiqi Liang*, Tianyang Liu*, Zilong Wang*, *et al.*

Blog Post, <https://cocoabench.github.io/>

Revisiting Reinforcement Learning for LLM Reasoning from A Cross-Domain Perspective

Zhoujun Cheng*, **Shibo Hao***, Tianyang Liu*, *et al.*

NeurIPS 2025 (Datasets & Benchmarks Track)

Reasoning by Superposition: A Theoretical Perspective on Chain of Continuous Thought

Hanlin Zhu*, **Shibo Hao***, Zhiting Hu, Jiantao Jiao, Stuart Russell, Yuandong Tian

NeurIPS 2025

K2-V2: A 360-Open, Reasoning-Enhanced LLM

Zhengzhong Liu, Liping Tang, Linghao Jin, Haonan Li, Nikhil Ranjan, Desai Fan, Shaurya Rohatgi, Richard Fan, Omkar Pangarkar, Huijuan Wang, Zhoujun Cheng, Suqi Sun, Seungwook Han, Bowen Tan, Gurpreet Gosal, Xudong Han, Varad Pimpalkhute, **Shibo Hao**, Ming Shan Hee, Joel Hestness, Haolong Jia, Liqun Ma, Aaryamonvikram Singh, Daria Soboleva, Natalia Vassilieva, Renxi Wang, Yingquan Wu, Yuekai Sun, Taylor Killian, Alexander Moreno, John Maggs, Hector Ren, Guowei He, Hongyi Wang, Xuezhe Ma, Yuqi Wang, Mikhail Yurochkin, Eric P. Xing

Technical Report, 2025

K2-Think: A Parameter-Efficient Reasoning System

Zhoujun Cheng*, Richard Fan*, **Shibo Hao***, Taylor W. Killian*, Haonan Li*, Suqi Sun*, *et al.*

Technical Report, 2025

Pandora: Towards General World Model with Natural Language Actions and Video States

Jiannan Xiang*, Guangyi Liu*, Yi Gu*, Qiyue Gao, Yuting Ning, Yuheng Zha, Zeyu Feng, Tianhua Tao, **Shibo Hao**, Yemin Shi, Zhengzhong Liu, Eric P. Xing, Zhiting Hu

Technical Report, 2024

Flow of Reasoning: Efficient Training of LLM Policy with Divergent Thinking

Fangxu Yu, Lai Jiang, Haoqiang Kang, **Shibo Hao**, Lianhui Qin

ICML 2025

LLM Pretraining with Continuous Concepts

Jihoon Tack, Jack Lanchantin, Jane Yu, Andrew Cohen, Ilia Kulikov, Janice Lan, **Shibo Hao**, Yuandong Tian, Jason Weston, Xian Li

ICLR 2026

Emergence of Superposition: Unveiling the Training Dynamics of Chain of Continuous Thought

Hanlin Zhu, **Shibo Hao**, Zhiting Hu, Jiantao Jiao, Stuart Russell, Yuandong Tian

ICLR 2026

Technical Skills

Programming: Python, C++, HTML, JavaScript

Deep Learning: Pytorch, Transformers, Distributed Training (e.g., FSDP), Accelerated Inference (SGLang/vLLM)

Languages: English – Fluent, Chinese – Native

Services

Reviewer for ICML (2024 - 2025), NeurIPS (2024), ICLR (2025), and ACL-ARR (Oct. 2023 - Now)