Shibo Hao

📞 858-789-3662 ✉ s5hao@ucsd.edu 🌐 ber666.github.io ⭘ github.com/ber666

Education

University of California, San Diego — Sep. 2022 – Present
Ph.D. in Data Science. Advisor: Zhiting Hu — La Jolla, CA

Peking University — Sep. 2018 – Jun. 2022
B.S. in Computer Science — Beijing, China

Experience

Meta FAIR Lab — Jun. 2024 – Jan. 2025
Research Scientist Intern. Advisor: Yuandong Tian, Jason Weston — Menlo Park, CA

Research Interests

My research goal is to push the boundaries of machine reasoning. My work includes building a system-2 reasoning framework using world-model planning (Pandora, RAP), training LLMs to reason with reinforcement learning (FoR, OREO), augmenting LLMs with external tools (ToolkenGPT) and exploring reasoning in latent space (COCONUT). We are maintaining a library (LLM Reasoners) for reasoning algorithms in a unified formulation.

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

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
arXiv preprint arXiv:2412.06769, 2024
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)
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
2.1k Stars (as of Mar. 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
ICLR 2025 Workshop on Reasoning and Planning for LLMs (Oral)

BertNet: Harvesting Knowledge Graphs with Arbitrary Relations from Pretrained Language Models
Shibo Hao*, Bowen Tan*, Kaiwen Tang, Bin Ni, Xiyan Shao, Hengzhe Zhang, Eric Xing, Zhiting Hu
Findings of ACL 2023

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
arXiv preprint arXiv:2406.09455, 2024

Flow of Reasoning: Efficient Training of LLM Policy with Divergent Thinking
Fangxu Yu, Lai Jiang, Haoqiang Kang, **Shibo Hao**, Lianhui Qin
arXiv preprint arXiv:2406.05673, 2024

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
arXiv preprint arXiv:2502.08524, 2025

Linear Correlation in LM's Compositional Generalization and Hallucination
Letian Peng, Chenyang An, **Shibo Hao**, Chengyu Dong, Jingbo Shang
arXiv preprint arXiv:2502.04520, 2025

Understanding the Sources of Uncertainty for Large Language and Multimodal Models
Ziran Yang, **Shibo Hao**, Hao Sun, Lai Jiang, Qiyue Gao, Yian Ma, Zhiting Hu
ICLR Workshop: Quantify Uncertainty and Hallucination in Foundation Models, 2025

Neural-symbolic Interaction and Co-evolving
Bowen Tan, **Shibo Hao**, Eric Xing, Zhiting Hu
Compendium of Neurosymbolic Artificial Intelligence 369, 125, 2023

Does Recommend-Revise Produce Reliable Annotations? An Analysis on Missing Instances in DocRED
Qian Huang, **Shibo Hao**, Yifan Ye, Shuang Zhu, Yansong Feng, Dongyan Zhao
ACL 2022

Benchmarking Commonsense Knowledge Base Population with an Effective Evaluation Dataset
Tianqing Fang*, Weiqi Wang*, Sehyun Choi, **Shibo Hao**, Hongming Zhang, Yangqiu Song, Ben He
EMNLP 2021

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

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