Haowei Lin

Beijing, China | linhaowei@pku.edu.cn | linhaowei1.github.io | github.com/linhaowei1

Research Overview

I hope to build next-generation AI with high real-world impact. My current work is centered on two primary areas: **AI Scientist**, where I design self-evolving code agents for scientific problems, and **Generative Foundation Models**, where I develop autoregressive and diffusion models across multi-modalities (e.g., image, text, molecules, MDP, 3D). My work has been published in ∼20 papers at top-tier ML/NLP conferences and journals, including a co-fisrt **Nature Machine Intelligence cover paper** and **two co-first spotlight papers** (ICML/NeurIPS). I am an active contributor to prominent open-source agent projects (Terminal Bench, OpenHands, OpenEvolve, AlgoTune), which have collectively received over 60,000 GitHub stars.

Education

Institude for Artificial Intelligence, Peking University — Ph.D. Student in Machine Learning — Sept 2023 – Jul 2028 (expected)

- **Advisors:** Prof. Jianzhu Ma and Prof. Yitao Liang and Prof. Di He

Yuanpei College, Peking University — B.E. in Artificial Intelligence (Honor Track) — Aug 2019 – Jul 2023

- **GPA:** 3.85/4.00 (rank 1st in Major); Top 5% in Yuanpei College (12/305)
- **Bachelor Thesis:** *IneqGPT: A Neural-Symbolic Approach towards Automated Inequality Proving*
- **Research Advisor:** Prof. Bing Liu (Distinguished Professor of UIC)

Publications

(* denotes equal contribution. Highlighted papers are marked with ★ highlighted.)

LLM & Agent

Terminal-Bench: a Benchmark for AI Agents in Terminal Environments 2025
★ highlighted

Mike Merrill, Alex Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Yeon Shin, Thomas Walshe, E. Kelly Buchanan, Junhong Shen, Guanghao Ye, **Haowei Lin**, ··· (70 people), Ludwig Schmidt
Used by OpenAI, Google DeepMind, Anthropic, Meta, DeepSeek, etc. for LLM evaluation in their model release.

Efficient and Asymptotically Unbiased Constrained Decoding for Large Language Models 2025

Haotian Ye, Himanshu Jain, Chong You, Ananda Theertha Suresh, **Haowei Lin**, James Zou, Felix Yu.
In The Twenty-Eighth International Conference on Artificial Intelligence and Statistics (**AISTATS 2025**).

Generative Evaluation of Complex Reasoning in Large Language Models 2025

Haowei Lin*, Xiangyu Wang*, Ruilin Yan*, Baizhou Huang, Haotian Ye, Jianhua Zhu, Zihao Wang, James Zou, Jianzhu Ma, Yitao Liang.
arXiv preprint arXiv:2504.02810.

MCU: An Evaluation Framework for Open-Ended Game Agents ★ highlighted 2025

Xinyue Zheng*, **Haowei Lin***, Kaichen He, Zihao Wang, Qiang Fu, Haobo Fu, Zilong Zheng, Yitao Liang.
In The Forty-second International Conference on Machine Learning (**ICML 2025 Spotlight**).

Jarvis-1: Open-World Multi-Task Agents with Memory-Augmented Multimodal Language Models 2024

Zihao Wang, Shaofei Cai, Anji Liu, Yonggang Jin, Jinbing Hou, Bowei Zhang, **Haowei Lin**, Zhaofeng He, Zilong Zheng, Yaodong Yang, Xiaojian Ma, Yitao Liang.
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI).

RAT: Retrieval Augmented Thoughts Elicit Context-Aware Reasoning in Long-Horizon Generation 2024

Zihao Wang, Anji Liu, **Haowei Lin**, Jiaqi Li, Xiaojian Ma, Yitao Liang.
arXiv preprint arXiv:2403.05313.

OmniJarvis: Unified Vision-Language-Action Tokenization Enables Open-World Instruction Following Agents 2024

Zihao Wang, Shaofei Cai, Zhancun Mu, **Haowei Lin**, Ceyao Zhang, Xuejie Liu, Qing Li, Anji Liu, Xiaojian Shawn Ma, Yitao Liang.
In Advances in Neural Information Processing Systems 37 (**NeurIPS 2024**).

Selecting Large Language Model to Fine-tune via Rectified Scaling Law 2024
★ highlighted

Haowei Lin*, Baizhou Huang*, Haotian Ye*, Qinyu Chen, Zihao Wang, Sujian Li, Jianzhu Ma, Xiaojun Wan, James Zou, Yitao Liang.
In The Forty-first International Conference on Machine Learning (**ICML 2024**).

GROOT-2: Weakly Supervised Multi-Modal Instruction Following Agents 2024

Shaofei Cai, Bowei Zhang, Zihao Wang, **Haowei Lin**, Xiaojian Ma, Anji Liu, Yitao Liang.
In The Twelfth International Conference on Learning Representations (**ICLR 2024**).

Generative Foundation Model

TFG-Flow: Training-free Guidance in Multimodal Generative Flow ★ highlighted 2025

Haowei Lin*, Shanda Li*, Haotian Ye, Yiming Yang, Stefano Ermon, Yitao Liang, Jianzhu Ma.
In The Thirteenth International Conference on Learning Representations (**ICLR 2025**).

TFG: Unified Training-Free Guidance for Diffusion Models ★ highlighted 2024

Haotian Ye*, **Haowei Lin***, Jiaqi Han*, Minkai Xu, Sheng Liu, Yitao Liang, Jianzhu Ma, James Y. Zou, Stefano Ermon.
In Advances in Neural Information Processing Systems 37 (**NeurIPS 2024 Spotlight**).

Uni-3DAR: Unified 3D Generation and Understanding via Autoregression on Compressed Spatial Tokens ★ highlighted 2025

Shuqi Lu*, **Haowei Lin***, Lin Yao*, Zhifeng Gao, Xiaohong Ji, Linfeng Zhang, Guolin Ke.
arXiv preprint arXiv:2503.16278.

AI4Science

A Neural Symbolic Model for Space Physics ★ highlighted 2025

Jie Ying*, **Haowei Lin***, Chao Yue*, Yajie Chen, Chao Xiao, Quanqi Shi, Yitao Liang, Shing-Tung Yau, Yuan Zhou, Jianzhu Ma.
Nature Machine Intelligence (Cover Paper).

Integrating Protein Dynamics into Structure-Based Drug Design via Full-Atom Stochastic Flows 2025

Xiangxin Zhou, Yi Xiao, **Haowei Lin**, Xinheng He, Jiaqi Guan, Yang Wang, Qiang Liu, Feng Zhou, Liang Wang, Jianzhu Ma.
In The Thirteenth International Conference on Learning Representations (**ICLR 2025**).

A Data-Driven Group Retrosynthesis Planning Model Inspired by Neurosymbolic Programming 2025

Xuefeng Zhang, **Haowei Lin**, Muhan Zhang, Yuan Zhou, Jianzhu Ma.
Nature Communications.

Peptide Design through Binding Interface Mimicry 2025

Xiangzhe Kong, Rui Jiao, **Haowei Lin**, Ruihan Guo, Wenbing Huang, Wei-Ying Ma, Zihua Wang, Yang Liu, Jianzhu Ma.
Nature Biomedical Engineering.

Open-World Learning

Class Incremental Learning via Likelihood Ratio-Based Task Prediction 2024
★ highlighted

Haowei Lin, Yijia Shao, Weinan Qian, Ningxin Pan, Yiduo Guo, Bing Liu.
In The Twelfth International Conference on Learning Representations (**ICLR 2024**).

FLatS: Principled Out-of-Distribution Detection with Feature-Based Likelihood Ratio Score 2023

Haowei Lin and Yuntian Gu.
In The 2023 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2023**).

Continual Pre-Training of Language Models ★ highlighted 2023

Zixuan Ke*, Yijia Shao*, **Haowei Lin***, Tatsuya Konishi, Gyuhak Kim, Bing Liu.
In The Eleventh International Conference on Learning Representations (**ICLR 2023**).

Continual Training of Language Models for Few-Shot Learning 2022

Zixuan Ke, **Haowei Lin**, Yijia Shao, Hu Xu, Lei Shu, Bing Liu.
In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**).

Adapting a Language Model While Preserving Its General Knowledge 2022

Zixuan Ke, Yijia Shao, **Haowei Lin**, Hu Xu, Lei Shu, Bing Liu.
In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (**EMNLP 2022**).

CMG: A Class-Mixed Generation Approach to Out-of-Distribution Detection 2022

Mengyu Wang, Yijia Shao, **Haowei Lin**, and others.
In Joint European Conference on Machine Learning and Knowledge Discovery in Databases (**ECML**).

Professional Experience

Wizard Intelligence Learning Lab, WizardQuant — Researcher, Post-train Team 2025.04 – now

- Developing novel post-training paradigms for self-evolving agents targeting open scientific problems.

DP Technology — Research Intern 2024.12 – 2025.09

- Developed **Uni3DAR**, a unified autoregressive 3D generative model for both understanding and generation.

Helixon — Research Intern 2022.09 – 2023.09

- Worked on multi-conformation protein structure prediction built on AlphaFold.
- Designed peptide-based therapeutics via structure generation and optimization workflows.

Search Group, ByteDance — Summer Intern 2022.06 – 2022.09

- Designed and launched a set of intention analysis models for Douyin Search Engine.

Teaching Assistant 2021 – now

- Spring 2025 & 2024, Introduction to Artificial Intelligence, Peking University.
- Spring 2023, Mathematical Foundation for Artificial Intelligence, Peking University.
- Fall 2022 & 2023 & 2024 & 2025, Natural Language Processing with Deep Learning, Peking University.
- Fall 2022, Introduction to Theoretical Computer Science, Peking University.
- Fall 2021, A Preliminary Seminar on Artificial Intelligence, Peking University.

Conference Reviewer 2022 – now

- ICML: 2022 (*Outstanding Reviewer Award*), 2023, 2024, **2025**; NeurIPS: 2022, 2023, 2024, **2025**; ICLR: 2024, **2025**; ACL, EMNLP (ARR): 2024, **2025**; COLM: 2024, **2025**; CVPR: 2024, **2025**; AISTATS: **2025**.

Awards

- National Scholarship, 2023
- National Scholarship, 2025
- Outstanding Bachelor Thesis, Peking University

Skills, Interests & Others

Programming Languages: Python, C/C++, LaTeX, Bash, SQL, HTML, CSS, JavaScript

Frameworks & Tools: PyTorch, OpenHands, VeRL, Transformers, Django, Git, CUDA, Spark, Bootstrap4

Languages: Mandarin (native), English (proficient, TOEFL 105), Hokkien (native), Japanese (JLPT N4-level)

Interests: Singing (*PKU Top-10 singers in 2025*), Guitar, Piano, Painting