Haoxiang Wang

haoxiang096@gmail.com
haoxiang-wang.github.io
haoxiang-wang-071414ab

Professional Experience

2024-Present **NVIDIA Research**
Research Scientist on Generative AI
Manager: Dr. Ming-Yu Liu

2023 Summer **Apple AI/ML**
Research Intern on Machine Learning
Paper: *Merging Vision Foundation Models towards Semantic and Spatial Understanding*
Mentors: Dr. Hadi Pouransari & Dr. Mehrdad Farajtabar & Dr. Oncel Tuzel

2022 Summer **Amazon Web Services**
Research Intern on Quantum Machine Learning
Paper: *Predicting Properties of Quantum Systems with Conditional Generative Models*
Mentors: Dr. Cedric Yen-Yu Lin & Dr. Peter Komar

2021 Summer **Waymo LLC**, (formerly the Google self-driving car project)
Research Intern on Machine Learning
Project: *Improving Robustness of Perception Systems with Mixture-of-Experts (MoE)*
Mentor: Dr. Zhao Chen

Education

2019–2024 **University of Illinois, Urbana-Champaign**, *Urbana*, IL, USA
Ph.D. *Electrical & Computer Engineering*
Research Area: *Machine Learning*. Thesis Advisors: Dr. Han Zhao and Dr. Bo Li

2015–2019 **University of Illinois, Urbana-Champaign**, *Urbana*, IL, USA
B.S. with Highest Distinctions. Double Majors:
(1) *Statistics & Computer Science* (2) *Physics*

Research Interests

Large Language Models, Vision-Language Models, Video Generation Models

Pre-training, Reinforcement Learning, Reasoning, Reward Modeling

Machine Learning Skills

PyTorch, Large-Scale Distributed Training (up to 8K H100 GPUs)

Inference Optimization, Reinforcement Learning for LLMs

Publications

(* indicates equal contribution.)

Foundation Models (Large Language Models and Multi-Modal Models)

2025 **Cosmos World Foundation Model Platform for Physical AI**
NVIDIA (Core contributor)
ArXiv Preprint, 2025.
Leading contributor of Cosmos-Autoregressive (Autoregressive Video Generation Models)

2025 **Cosmos-Reason1: From Physical Common Sense To Embodied Reasoning**
NVIDIA (Core contributor)
ArXiv Preprint, 2025.
Core contributor of the SFT and RL of Cosmos-Reason1 (Vision-Language Model for Physical AI)

2025 **Bridging Supervised Learning and Reinforcement Learning in Math Reasoning**
Huayu Chen, Kaiwen Zheng, Qinsheng Zhang, Ganqu Cui, Yin Cui, Haotian Ye, Tsung-Yi Lin, Ming-Yu Liu, Jun Zhu, Haoxiang Wang
ArXiv Preprint, 2025.

EMNLP 2024 **Interpretable Preferences via Multi-Objective Reward Modeling and Mixture-of-Experts**
Haoxiang Wang*, Wei Xiong*, Tengyang Xie, Han Zhao, Tong Zhang
EMNLP Findings, 2024.

EMNLP 2024 **Semi-Supervised Reward Modeling via Iterative Self-Training**
Yifei He*, Haoxiang Wang*, Ziyan Jiang, Alexandros Papangelis, Han Zhao
EMNLP Findings, 2024.

TMLR 2024 **RLHF Workflow: From Reward Modeling to Online RLHF**
Hanze Dong*, Wei Xiong*, Bo Pang*, Haoxiang Wang*, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, Tong Zhang
Transactions on Machine Learning Research (**TMLR**), 2024

ACL 2024 **Arithmetic Control of LLMs for Diverse User Preferences: Directional Preference Alignment with Multi-Objective Rewards**
Haoxiang Wang*, Yong Lin*, Wei Xiong*, Rui Yang, Shizhe Diao, Shuang Qiu, Han Zhao, Tong Zhang
ACL, 2024.

TMLR 2024 **Enhancing Compositional Generalization via Compositional Feature Alignment**
Haoxiang Wang*, Haozhe Si*, Han Zhao
Transactions on Machine Learning Research (**TMLR**), 2024

CVPRW 2024 **SAM-CLIP: Merging Vision Foundation Models towards Semantic and Spatial Understanding**
Haoxiang Wang, Pavan Kumar Anasosalu Vasu, Fartash Faghri, Raviteja Vemulapalli, Mehrdad Farajtabar, Sachin Mehta, Mohammad Rastegari, Oncel Tuzel, Hadi Pouransari
CVPR Workshop, Efficient Large Vision Models (Spotlight), 2024.

Multi-Task Learning and Meta-Learning

2022 **Predicting Properties of Quantum Systems with Conditional Generative Models**
Haoxiang Wang*, Maurice Weber*, Josh Izaac, Cedric Yen-Yu Lin
Under review of **Quantum**

CVPR 2022 **Global Convergence of MAML and Theory-Inspired Neural Architecture Search for Few-Shot Learning**
Haoxiang Wang*, Yite Wang*, Ruoyu Sun, Bo Li
Conference on Computer Vision and Pattern Recognition (**CVPR**), 2022.

ICML 2021 **Bridging Multi-Task Learning and Meta-Learning: Towards Efficient Training and Effective Adaptation**
Haoxiang Wang, Han Zhao, Bo Li
International Conference on Machine Learning (**ICML**), 2021.

Robust and Trustworthy AI under Data Distribution Shifts

JMLR 2024 **Gradual Domain Adaptation: Theory and Algorithms**
Yifei He*, Haoxiang Wang*, Bo Li, Han Zhao
Journal of Machine Learning Research (**JMLR**), 2024

2023 **Invariant Feature Subspace Recovery: A New Class of Provable Domain Generalization Algorithms**
Haoxiang Wang*, Gargi Balasubramaniam*, Haozhe Si, Bo Li, Han Zhao
Under review of Journal of Machine Learning Research (**JMLR**)

ICML 2022 **Provable Domain Generalization via Invariant-Feature Subspace Recovery**
Haoxiang Wang, Haozhe Si, Bo Li, Han Zhao
International Conference on Machine Learning (**ICML**), 2022.

ICML 2022 **Understanding Gradual Domain Adaptation: Improved Analysis, Optimal Path and Beyond**
Haoxiang Wang, Bo Li, Han Zhao
International Conference on Machine Learning (**ICML**), 2022.

UAI 2022 **Future Gradient Descent for Adapting the Temporal Shifting Data Distribution in Online Recommendation System**
Mao Ye, Ruichen Jiang, Haoxiang Wang, Dhruv Choudhary, Xiaocong Du, Bhargav Bhushanam, Aryan Mokhtari, Arun Kejariwal, Qiang Liu
Conference on Uncertainty in Artificial Intelligence (**UAI**), 2022.

Optimization

NeurIPS 2019 **Learning Positive Functions with Pseudo Mirror Descent**
Yingxiang Yang, Haoxiang Wang, Negar Kiyavash, Niao He
Neural Information Processing Systems (**NeurIPS**), 2019. **(Spotlight Presentation)**

Open-Source Software

2024 **RLHFlow - Reward Modeling**, *https://github.com/RLHFlow/RLHF-Reward-Modeling*, Open-source code to train reward models for RLHF/alignment of large language models.

Service

2020-Now **Reviewer**, *ICML, NeurIPS, ICLR, ICCV, CVPR, TMLR, AISTATS, AAAI*