

HAO WU

Ph.D. Student | Research Intern @ Tencent IEG, LightSpeed Studios

Tsinghua University • Beijing, China

✉ wu-h25@mails.tsinghua.edu.cn | 📞 +86 182 2584 0735 | 📍 Beijing, Haidian

🌐 Homepage | 🎓 Google Scholar | 🏠 GitHub | 🏛️ Tsinghua University

I am a Ph.D. student at **Tsinghua University** and currently a Research Intern at **Tencent IEG, LightSpeed Studios** working on video generation. Previously, I conducted research at **Tencent CSIG, Tencent Jarvis Lab**; I graduated from the Department of Computer Science at the **University of Science and Technology of China (USTC)**, during which I was also a joint training student in the large model training group of the Machine Learning Platform Department at **Tencent Hunyuan**.

Research Interests: My research lies at the intersection of (1) **Robot Video Generation & World Models**; (2) **Multimodal Large Language Models (VLMs)**; (3) **Agentic Reasoning Systems**; and (4) **Large-scale Spatiotemporal Forecasting**. Broadly, I aim to build intelligent systems that can *understand, predict, and reason about* the physical world across space, time, and modalities. My work has appeared in top venues including **ICLR, NeurIPS, ICML, KDD, AAAI, ICCV, ACM MM, TKDE, and TPAMI**, with **1,200+ Google Scholar citations** and nearly **30 CCF-A publications** (including ~20 as first or corresponding author). I serve as a reviewer for **ICLR, NeurIPS, ICML, KDD, ICCV, AAAI, ACM MM, TPAMI, and TKDE**.

RESEARCH EXPERIENCE

- | | |
|--|----------------------|
| Research Intern, Tencent IEG, LightSpeed Studios | May 2026 – Present |
| • Researching video generation world models , with a focus on long-horizon video generation, controllable dynamics, and physically grounded simulation for interactive scenarios. | |
| Research Intern, Tencent CSIG, Tencent Jarvis Lab | Aug 2025 – May 2026 |
| • Continued research on multimodal foundation models, agent systems, and world models in industrial-scale settings. | |
| Research Intern, Tencent Hunyuan, Machine Learning Platform Department | Aug 2023 – Jul. 2025 |
| • Worked on large models, world models, and multimodal generative modeling in Tencent Hunyuan. | |
| Online Research Intern, UCLA | May 2023 – May 2024 |
| • Advisor: Xiao Luo . Remote research on multimodal learning and dynamics modeling. | |
| Research Intern, CityMind Lab, HKUST (Guangzhou) | Mar 2023 – Aug 2023 |
| • Advisors: Yuxuan Liang and Kun Wang . Early-stage research on spatiotemporal modeling. | |

SELECTED AWARDS & HONORS

- | | |
|--|------|
| Outstanding Graduate, USTC × Tencent Joint Training Program | 2025 |
| National Scholarship (Top 1% in China), USTC | 2022 |
| First-Class Scholarship, USTC | 2022 |

PUBLICATIONS

† equal contribution, ^ co-first author, # / * corresponding author. Nearly 30 CCF-A publications; full list on Google Scholar.

Robot Video Generation & World Models

RoboAlign-R1: Distilled Multimodal Reward Alignment for Robot Video World Models *Under Review*, 2026

Hao Wu, Yuqi Li, Yuan Gao, Fan Xu, Fan Zhang, Kun Wang, Penghao Zhao, Qiufeng Wang, Yizhou Zhao, Weiyan Wang, Yingli Tian, Xian Wu, Xiaomeng Huang#

Under Review. [Project Page]

Multimodal Large Language Models & Agent Systems

SEAT: Calibrating Inference Time Alignment with Sequence-level Risk Accumulation

ACL, 2026

Shanwen Tan[†], Ziyang Dong[†], Wei Ju[#], Yiwei Fu[#], **Hao Wu**, Kun Wang, Yifan Wang, Ziyue Qiao
Annual Meeting of the Association for Computational Linguistics (ACL), 2026

MME-Emotion: A Holistic Evaluation Benchmark for Emotional Intelligence in Multimodal Large Language Models

ICLR, 2026

Fan Zhang, Zebang Cheng, Chong Deng, Haoxuan Li, Zheng Lian, Qian Chen, Huadai Liu, Wen Wang, YiFan Zhang, Renrui Zhang, Ziyu Guo, Zhihong Zhu, **Hao Wu**, Haixin Wang, Yefeng Zheng, Xiaojiang Peng, Xian Wu, Kun Wang, Xiangang Li, Jieping Ye, Pheng-Ann Heng

International Conference on Learning Representations (ICLR), 2026 (Core Rank A)

Safety-Aware Rollouts with Self-Reflection and Structured Rewards

TPAMI, Under Review

Huahui Yi, Kun Wang, Haolong Hu, Moayad Aloqaily, Liang Lin, Junhao Dong, Qiankun Li, Xing Fan, **Hao Wu**[#], Yang Liu, Qingsong Wen

IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Under Review

Large-scale Spatiotemporal & Earth System Forecasting

TritonCast: Advanced Long-term Earth System Forecasting

NMI, Under Revision

Hao Wu, Yuan Gao, Ruijian Gou, Xian Wu, Chuhan Wu, Huahui Yi, Johannes Brandstetter, Fan Xu, Niklas Boers, Kun Wang, Penghao Zhao, Hao Jia, Qi Song, Xinliang Liu, Juncai He, Shuhao Cao, Huanshuo Dong, Yanfei Xiang, Fan Zhang, Haixin Wang, Xingjian Shi, Qiufeng Wang, Shuaipeng Li, Ruobing Xie, Feng Tao, Yuxu Lu, Yu Guo, Yuntian Chen, Yuxuan Liang, Qingsong Wen, Wanli Ouyang, Deliang Chen, Xiaomeng Huang

Nature Machine Intelligence, In Peer Review

PnP-Corrector: A Universal Correction Framework for Coupled Spatiotemporal Forecasting

ICML, 2026

Hao Wu, Fan Xu, Yuxu Lu, Penghao Zhao, Fan Zhang, Hao Jia, Yuxuan Liang, Ruijian Gou, Qingsong Wen, Xian Wu, Xiaomeng Huang, Yuan Gao

International Conference on Machine Learning (ICML), 2026

NuwaDynamics+: A Causality-Aware Generative Framework for Spatio-Temporal Representation Learning

IEEE TPAMI, 2026

Kun Wang, Yifan Duan, **Hao Wu**, Jian Zhao, Kai Wang, Zhengyang Zhou, Yuxuan Liang, Xu Wang^{*}, Yang Wang^{*}, Yu Zheng, Xuelong Li

IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2026

NeuralOM: Neural Ocean Model for Subseasonal-to-Seasonal Simulation

AAAI, 2026

Yuan Gao[†], **Hao Wu**^{†‡}, Fan Xu, Yanfei Xiang, Ruijian Gou, Ruiqi Shu, Qingsong Wen, Xian Wu, Kun Wang^{*}, Xiaomeng Huang^{*}

AAAI Conference on Artificial Intelligence, 2026

OneForecast: A Universal Framework for Global and Regional Weather Forecasting

ICML, 2025

Yuan Gao, **Hao Wu**, Ruiqi Shu, Huanshuo Dong, Fan Xu, Rui Ray Chen, Yibo Yan, Qingsong Wen, Xuming Hu, Kun Wang, Jiahao Wu, Li Qing, Hui Xiong, Xiaomeng Huang[#]

International Conference on Machine Learning (ICML), 2025

Frequency-Aligned Knowledge Distillation for Lightweight Spatiotemporal Forecasting

ICCV, 2025

Yuqi Li, Chuanguang Yang, Hansheng Zeng, Zeyu Dong, Zhulin An, Yongjun Xu, Yingli Tian, **Hao Wu**[#]

IEEE/CVF International Conference on Computer Vision (ICCV), 2025

DynST: Dynamic Sparse Training for Resource-Constrained Spatio-Temporal Forecasting *KDD*, 2025

Hao Wu, Haomin Wen, Guibin Zhang, Yutong Xia, Yuxuan Liang, Yu Zheng, Qingsong Wen, Kun Wang#
ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2025

NuwaDynamics: Discovering and Updating in Causal Spatio-Temporal Modeling *ICLR*, 2024

Kun Wang[^], *Hao Wu*[^], Yifan Duan, Guibin Zhang, Kai Wang, Xiaojiang Peng, Yu Zheng, Yuxuan Liang#, Yang Wang#
International Conference on Learning Representations (ICLR), 2024

EarthFarseer: Versatile Spatio-Temporal Dynamical Systems Modeling in One Model *AAAI*, 2024

Hao Wu, Yuxuan Liang, Wei Xiong#, Zhengyang Zhou, Wei Huang, Shilong Wang, Kun Wang#
AAAI Conference on Artificial Intelligence, 2024

Modeling Spatio-Temporal Dynamical Systems with Neural Discrete Learning and Levels-of-Experts *TKDE*, 2024

Kun Wang[^], *Hao Wu*[^], Guibin Zhang, Junfeng Fang, Yuxuan Liang, Yuankai Wu, Roger Zimmermann, Yang Wang#
IEEE Transactions on Knowledge and Data Engineering (TKDE), 2024

PastNet: Introducing Physical Inductive Biases for Spatio-Temporal Video Prediction *ACM MM*, 2024

Hao Wu, Wei Xiong, Fan Xu, Xiao Luo#, Chong Chen, Xian-Sheng Hua, Haixin Wang#
ACM International Conference on Multimedia (ACM MM), 2024

Scientific Machine Learning & Neural Operators

Differential-Integral Neural Operator for Long-Term Turbulence Forecasting *KDD*, 2026

Hao Wu, Yuan Gao, Fan Xu, Fan Zhang, Qingsong Wen, Xiaomeng Huang, Xian Wu
ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2026 (Research Track, Accepted)

Advanced Global Wildfire Activity Modeling with Hierarchical Graph ODE *KDD*, 2026

Fan Xu, Wei Gong, *Hao Wu*, Lilan Peng, Nan Wang, Qingsong Wen, Xian Wu, Kun Wang, Xibin Zhao
ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2026

Breaking the Discretization Barrier of Continuous Physics Simulation Learning *NeurIPS*, 2025

Fan Xu, *Hao Wu*, Nan Wang, Lilan Peng, Kun Wang, Wei Gong, Xibin Zhao
Advances in Neural Information Processing Systems (NeurIPS), 2025

PURE: Prompt Evolution with Graph ODE for Out-of-distribution Fluid Dynamics Modeling *NeurIPS*, 2024

Hao Wu, Changhu Wang, Fan Xu, Jinbao Xue, Chong Chen, Xian-Sheng Hua, Xiao Luo#
Advances in Neural Information Processing Systems (NeurIPS), 2024

Prometheus: Out-of-distribution Fluid Dynamics Modeling with Disentangled Graph ODE *ICML*, 2024

Hao Wu, Huiyuan Wang, Kun Wang, Weiyan Wang, Changanye Ye, Yangyu Tao, Chong Chen, Xian-Sheng Hua, Xiao Luo#

International Conference on Machine Learning (ICML), 2024

Neural Manifold Operators for Learning the Evolution of Physical Dynamics *KDD*, 2024

Hao Wu, Kangyu Weng, Shuyi Zhou, Xiaomeng Huang#, Wei Xiong#
ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024

IDEA: An Invariant Perspective for Efficient Domain Adaptive Image Retrieval *NeurIPS*, 2023

Haixin Wang[^], *Hao Wu*[^], Jinan Sun, Shikun Zhang, Chong Chen, Xian-Sheng Hua, Xiao Luo#

Advances in Neural Information Processing Systems (NeurIPS), 2023

ACADEMIC SERVICE

Reviewer: ICLR, NeurIPS, ICML, KDD, ICCV, AAAI, ACM MM, TPAMI, TKDE

TECHNICAL SKILLS

Programming Languages: Python, C++, C, Java

Frameworks & Tools: PyTorch, JAX, TensorFlow, DeepSpeed, Megatron, Git, Linux, \LaTeX

Research Areas: Robot Video Generation & World Models, Multimodal LLMs / VLMs, Agent Systems, Reinforcement Learning from Human / AI Feedback, Spatiotemporal & Earth-System Forecasting, Neural Operators