Haoxiang Ren

renhx@lamda.nju.edu.cn / rhx11111.github.io / Google Scholar

PROFILE

I am a master's student in the LAMDA Group at Nanjing University, advised by Prof. Yang Yu.

EDUCATION

Master's Student in Artificial Intelligence 2026 - Present
Nanjing University, Nanjing, China
LAMDA Group, School of Artificial Intelligence

Bachelor of Artificial Intelligence 2022 - 2026
Nanjing University, Nanjing, China
School of Artificial Intelligence

RESEARCH INTERESTS

My research aims to develop general-purpose embodied intelligence through reinforcement learning. I focus on generative world models and model-based reinforcement learning for policy learning, planning, and adaptation in embodied environments.

REPRESENTATIVE PUBLICATIONS

Towards Practical World Model-based Reinforcement Learning for Vision-Language-Action Models
ICML, 2026
*Co-first author; **Got OpenAI's interest!***

- Proposed VLA-MBPO to fine-tune VLA policies in learned world models, combining unified multimodal prediction of observations and rewards, coordinated decoding across camera views, and short branched rollouts over action chunks.
- Demonstrated gains in policy performance and sample efficiency across LIBERO and real-world tasks on bimanual and whole-body robots.

When to Trust Your Teacher: Uncertainty-Aware On-Policy Distillation for Embodied Policies
RSS Workshop on Diffusion for Robot Learning, 2026
Co-first author

- Proposed UA-OPD for reliable on-policy distillation of embodied policies, combining self-consistency estimates of teacher uncertainty with uncertainty penalties and trajectory truncation to reduce unreliable supervision on student-visited states.
- Demonstrated faster convergence and stronger performance than standard on-policy distillation on LIBERO and RoboTwin 2.0, with student policies outperforming the teacher on most evaluated tasks.

Prioritized Rollouts for Efficient World Model-based Vision-Language-Action Policy Optimization
arXiv preprint, 2026
Co-first author

- Proposed U-GROW, a lightweight sampling layer for model-based VLA training that filters failure tails and prioritizes informative, high-uncertainty states for rollout generation without changing the policy optimization objective.
- Achieved better sample efficiency and policy performance on RoboTwin 2.0, LIBERO, and real-world manipulation using the same world model and rollout budget as VLA-MBPO.

FULL PUBLICATIONS

Prioritized Rollouts for Efficient World Model-based Vision-Language-Action Policy Optimization
Yifei Sheng*, **Haoxiang Ren***, Zhilong Zhang*, Haonan Wang*, Runjie Xu, Yihao Sun, Nan Tang, Zhichao Wu, Lei Yuan, Haoxin Lin, Yang Yu
arXiv preprint, 2026

When to Trust Your Teacher: Uncertainty-Aware On-Policy Distillation for Embodied Policies
Haoxiang Ren*, Yifei Sheng*, Zhilong Zhang*, Runjie Xu, Bowen Lu, Haoxin Lin, Yang Yu
RSS Workshop on Diffusion for Robot Learning, 2026

Towards Practical World Model-based Reinforcement Learning for Vision-Language-Action Models
Zhilong Zhang*, **Haoxiang Ren***, Yihao Sun*, Yifei Sheng, Haonan Wang, Haoxin Lin, Zhichao Wu, Pierre-Luc Bacon, Yang Yu
ICML, 2026

Speedup Patch: Learning a Plug-and-Play Policy to Accelerate Embodied Manipulation
Zhichao Wu*, Junyin Ye*, Zhilong Zhang*, Yihao Sun, Haoxin Lin, Jiaheng Luo, **Haoxiang Ren**, Lei Yuan, Yang Yu
ICML, 2026

InCLET: In-Context Learning Embeddings from LLM-Imagined Trajectories for Embodied Robot Instruction Following
Peng-Yuan Wang, Chen-Yang Wang, Jing-Cheng Pang, **Haoxiang Ren**, Zhilong Zhang, Xuhui Liu, Tian-Shuo Liu, Hong Qian, Yang Yu
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2026

Anticipation-VLA: Solving Long-Horizon Embodied Tasks via Anticipation-based Subgoal Generation
Zhilong Zhang*, Wenyu Luo*, Haonan Wang*, Yifei Sheng*, Yidi Wang, Hanyuan Guo, **Haoxiang Ren**, Xinghao Du, Yuhan Che, Tongtong Cao, Lei Yuan, Yang Yu
arXiv preprint, 2026

How Should World Models Be Evaluated for Embodied Decision-Making? A Decision-Making-Centric Position
Yang Yu, Shiyuan Zhang, Yifei Sheng, **Haoxiang Ren**, Haoxin Lin
arXiv preprint, 2026

On the Representation Degradation in Vision-Language-Action Models
Zhilong Zhang, Xiong-Hui Chen, Yidi Wang, Yihao Sun, Wenyu Luo, **Haoxiang Ren**, Haoxin Lin, Yang Yu
OpenReview manuscript, 2026

WHALE: Towards Generalizable and Scalable World Models for Embodied Decision-making
Zhilong Zhang*, Ruifeng Chen*, Junyin Ye*, Yihao Sun, **Haoxiang Ren**, Xinghao Du, Pengyuan Wang, Jing-Cheng Pang, Kaiyuan Li, Tian-Shuo Liu, Haoxin Lin, Yang Yu, Zhi-Hua Zhou
NeurIPS Workshop on Embodied World Models for Decision Making, 2025

** Equal contribution.*