

# ZHENFEI(JEREMY) YIN

✉ [jeremyyin@robots.ox.ac.uk](mailto:jeremyyin@robots.ox.ac.uk) 🏠 <https://yinzhenfei.github.io>

## RESEARCH INTERESTS

Agentic AI • Robotics • AI Scientists

## EDUCATION AND WORKING EXPERIENCE

- 2026.03-Present Postdoctoral Researcher in the Department of Engineering Science at the University of Oxford  
Research Highlight: General AI Scientists empowered by LLM-Agents and Robotics  
Advisor: Philip Torr
- 2025.04-Present Visiting Research Fellow in the Department of Engineering Science at the University of Oxford  
Research Highlight: AI Scientists based on Foundation Models  
Advisor: Philip Torr
- 2022.07-2025.12 Ph.D. in the School of Computer Science at the University of Sydney  
Thesis Title: Exploring Multi-Modal Representation Learning for Embodied Agents.  
Advisor: Wanli Ouyang and Zhiyong Wang
- 2022.07-2025.12 Rising Star Research Fellow at the Shanghai AI Lab  
Working on AI Agents  
Mentor: Lei Bai, Amanda Shao
- 2018.06-2022.07 Research Manager at the AGI Group of SenseTime  
Working on Multi-Modal Foundation Models  
Mentor: Junjie Yan
- 2015.09-2019.06 B.Sc. in the School of Software Engineering at HUST  
ACM/ICPC Team Captain, World Finalist  
Advisor: Kun Hu

## PUBLICATIONS & PREPRINTS

Equal contributions are denoted using \*, while corresponding authors are indicated using ‡. Project leads are marked with †, which indicates co-advising role without being the corresponding author. As research on foundation models and AI agents often involves systematic collaborative efforts, resulting in some technical reports using alphabetical author ordering. When available, DOIs and links to preprints are provided. Publications are listed in reverse order of their first arXiv submission.

- 2025 Yuyang Hu\*, Shichun Liu\*, Yanwei Yue\*, Guibin Zhang\*, Boyang Liu, Fangyi Zhu, Jiahang Lin, Honglin Guo, Shihan Dou, Zhiheng Xi, Senjie Jin, Jiejun Tan, Yanbin Yin, Jiongnan Liu, Zeyu Zhang, Zhongxiang Sun, Yutao Zhu, Hao Sun, Boci Peng, Zhenrong Cheng, Xuanbo Fan, Jiaxin Guo, Xinlei Yu, Zhenhong Zhou, Zewen Hu, Jiahao Huo, Junhao Wang, Yuwei Niu, Yu Wang, Zhenfei Yin, Xiaobin Hu, Yue Liao, Qiankun Li, Kun Wang, Wangchunshu Zhou, Yixin Liu, Dawei Cheng, Qi Zhang, Tao Gui, Shirui Pan, Yan Zhang, Philip Torr, Zhicheng Dou, Ji-Rong Wen, Xuanjing Huang, Yu-Gang Jiang, Shuicheng Yan. Memory in the Age of AI Agents. Preprint. arXiv: [2512.13564](https://arxiv.org/abs/2512.13564)

Zonghao Ying\*, Yangguang Shao\*, Jianle Gan, Gan Xu, Junjie Shen, Wenxin Zhang, Quanchen Zou, Junzheng Shi, [Zhenfei Yin](#), Mingchuan Zhang, Aishan Liu, Xianglong Liu. SecureWebArena: A Holistic Security Evaluation Benchmark for LVLm-based Web Agents. Preprint. arXiv: [2510.10073 \[cs.AI\]](#)<sup>↗</sup>

Xiangyuan Xue, Yifan Zhou, Guibin Zhang, Zaibin Zhang, Yijiang Li, Chen Zhang, [Zhenfei Yin](#)<sup>‡</sup>, Philip Torr, Wanli Ouyang, Lei Bai<sup>‡</sup>. CoMAS: Co-Evolving Multi-Agent Systems via Interaction Rewards. Preprint. arXiv: [2510.08529 \[cs.AI\]](#)<sup>↗</sup>

Qianshan Wei\*, Tengchao Yang\*, Yaochen Wang\*, Xinfeng Li<sup>‡</sup>, Lijun Li, [Zhenfei Yin](#), Yi Zhan, Thorsten Holz, Zhiqiang Lin, XiaoFeng Wang. A-MemGuard: A Proactive Defense Framework for LLM-Based Agent Memory. Preprint. arXiv: [2510.02373 \[cs.AI\]](#)<sup>↗</sup>

Guibin Zhang, Fanci Meng, Guancheng Wan, Zherui Li, Kun Wang, [Zhenfei Yin](#), Lei Bai, Shuicheng Yan. LatentEvolve: Self-Evolving Test-Time Scaling in Latent Space. Preprint. arXiv: [2509.24771 \[cs.AI\]](#)<sup>↗</sup>

Zelin Tan, Hejia Geng, Mulei Zhang, Xiaohang Yu, Guancheng Wan, Yifan Zhou, Qiang He, Xiangyuan Xue, Heng Zhou, Yutao Fan, Zhongzhi Li, Zaibin Zhang, Guibin Zhang, Chen Zhang<sup>‡</sup>, [Zhenfei Yin](#)<sup>‡</sup>, Lei Bai. Scaling Behaviors of LLM Reinforcement Learning Post-Training: An Empirical Study in Mathematical Reasoning. Preprint. arXiv: [2509.25300 \[cs.AI\]](#)<sup>↗</sup>

Guancheng Wan\*, Leixin Sun\*, Longxu Dou, Zitong Shi, Fang Wu, Eric Hanchen Jiang, Wenke Huang, Guibin Zhang, Hejia Geng, Xiangru Tang, [Zhenfei Yin](#)<sup>‡</sup>, Yizhou Sun, Wei Wang. Diagnose, Localize, Align: A Full-Stack Framework for Reliable LLM Multi-Agent Systems under Instruction Conflicts. Preprint. arXiv: [2509.23188 \[cs.AI\]](#)<sup>↗</sup>

Xiangru Tang\*, Wanghan Xu\*, Yujie Wang\*, Zijie Guo\*, Daniel Shao, Jiapeng Chen, Cixuan Zhang, Ziyi Wang, Lixin Zhang, Guancheng Wan, Wenlong Zhang, Lei Bai, [Zhenfei Yin](#)<sup>‡</sup>, Philip Torr, Hanrui Wang, Di Jin. Eigen-1: Adaptive Multi-Agent Refinement with Monitor-Based RAG for Scientific Reasoning. Preprint. arXiv: [2509.21193 \[cs.AI\]](#)<sup>↗</sup>

Wenxuan Huang, Shuang Chen, Zheyong Xie, Shaosheng Cao<sup>‡</sup>, Shixiang Tang, Yufan Shen, Qingyu Yin, Wenbo Hu, Xiaoman Wang, Yuntian Tang, Junbo Qiao, Yue Guo, Yao Hu, [Zhenfei Yin](#)<sup>‡</sup>, Philip Torr, Yu Cheng, Wanli Ouyang, Shaohui Lin<sup>‡</sup>. Interleaving Reasoning for Better Text-to-Image Generation. Preprint. arXiv: [2509.06945 \[cs.CV\]](#)<sup>↗</sup>

Guibin Zhang\*, Hejia Geng\*, Xiaohang Yu\*, [Zhenfei Yin](#)<sup>‡</sup>, Zaibin Zhang, Zelin Tan, Heng Zhou, Zhongzhi Li, Xiangyuan Xue, Yijiang Li, Yifan Zhou, Yang Chen, Chen Zhang, Yutao Fan, Zihu Wang, Songtao Huang, Yue Liao, Hongru Wang, Mengyue Yang, Heng Ji, Michael Littman, Jun Wang, Shuicheng Yan, Philip Torr, Lei Bai<sup>‡</sup>. The Landscape of Agentic Reinforcement Learning for LLMs: A Survey. Preprint. arXiv: [2509.02547 \[cs.AI\]](#)<sup>↗</sup>

Zhiheng Xi, Guanyu Li, Yutao Fan, Honglin Guo, Yufang Liu, Xiaoran Fan, Jiaqi Liu, Jingchao Ding, Wangmeng Zuo, [Zhenfei Yin](#)<sup>‡</sup>, Lei Bai, Tao Ji, Tao Gui<sup>‡</sup>, Qi Zhang, Philip Torr, Xuanjing Huang. BMMR: A Large-Scale Bilingual Multimodal Multi-Discipline Reasoning Dataset. Neural Information Processing Systems, NeurIPS 2025, Datasets and Benchmarks Track.

Li Kang, Xiufeng Song, Heng Zhou\*, Yiran Qin<sup>‡</sup>, Jie Yang, Xiaohong Liu, Philip Torr, Lei Bai<sup>‡</sup>, [Zhenfei Yin](#)<sup>‡</sup>. VIKI-R: Coordinating Embodied Multi-Agent Cooperation via Reinforcement Learning. Neural Information Processing Systems, NeurIPS 2025. arXiv: [2506.09049 \[cs.AI\]](#)<sup>↗</sup>

Rui Li, Zixuan Hu, Wenxi Qu\*, Jinouwen Zhang, [Zhenfei Yin](#), Sha Zhang, Xuantuo Huang, Hanqing Wang, Tai Wang, Jiangmiao Pang, Wanli Ouyang, Lei Bai, Wangmeng Zuo, Ling-Yu Duan, Dongzhan Zhou<sup>‡</sup>, Shixiang Tang<sup>‡</sup>. LabUtopia: High-Fidelity Simulation and Hierarchical Benchmark for Scientific Embodied Agents. Neural Information Processing Systems, NeurIPS 2025.

Guibin Zhang\*, Hejia Geng\*, Xiaohang Yu\*, Zhenfei Yin<sup>‡</sup>, Zaibin Zhang, Zelin Tan, Heng Zhou, Zhongzhi Li, Xiangyuan Xue, Yijiang Li, Yifan Zhou, Yang Chen, Chen Zhang, Yutao Fan, Zihu Wang, Songtao Huang, Yue Liao, Hongru Wang, Mengyue Yang, Heng Ji, Michael Littman, Jun Wang, Shuicheng Yan, Philip Torr, Lei Bai<sup>‡</sup>. The Landscape of Agentic Reinforcement Learning for LLMs: A Survey. Preprint. arXiv: 2509.02547 [cs.AI]<sup>☞</sup>

Xiangyuan Xue, Yifan Zhou, Guibin Zhang, Zaibin Zhang, Yijiang Li, Chen Zhang, Zhenfei Yin<sup>‡</sup>, Philip Torr, Wanli Ouyang, Lei Bai<sup>‡</sup>. CoMAS: Co-Evolving Multi-Agent Systems via Interaction Rewards. Preprint. arXiv: [cs.AI]<sup>☞</sup>

Qianshan Wei, Tengchao Yang, Yaochen Wang\*, Xinfeng Li<sup>‡</sup>, Lijun Li, Zhenfei Yin, Yi Zhan, Thorsten Holz, Zhiqiang Lin, XiaoFeng Wang. A-MemGuard: A Proactive Defense Framework for LLM-Based Agent Memory. Preprint.

Guibin Zhang, Fanci Meng, Guancheng Wan, Zherui Li, Kun Wang, Zhenfei Yin, Lei Bai, Shuicheng Yan. LatentEvolve: Self-Evolving Test-Time Scaling in Latent Space. Preprint.

Zelin Tan, Hejia Geng, Mulei Zhang, Xiaohang Yu, Guancheng Wan, Yifan Zhou, Qiang He, Xiangyuan Xue, Heng Zhou, Yutao Fan, Zhongzhi Li, Zaibin Zhang, Guibin Zhang, Chen Zhang<sup>‡</sup>, Zhenfei Yin<sup>‡</sup>, Lei Bai. Scaling Behaviors of LLM Reinforcement Learning Post-Training: An Empirical Study in Mathematical Reasoning. Preprint.

Guancheng Wan, Leixin Sun, Longxu Dou, Zitong Shi, Fang Wu, Eric Hanchen Jiang, Wenke Huang, Guibin Zhang, Hejia Geng, Xiangru Tang, Zhenfei Yin<sup>‡</sup>, Yizhou Sun, Wei Wang. Diagnose, Localize, Align: A Full-Stack Framework for Reliable LLM Multi-Agent Systems under Instruction Conflicts. Preprint.

Xiangru Tang, Wanghan Xu, Yujie Wang, Zijie Guo, Daniel Shao, Jiapeng Chen, Cixuan Zhang, Ziyi Wang, Lixin Zhang, Guancheng Wan, Wenlong Zhang, Lei Bai, Zhenfei Yin<sup>‡</sup>, Philip Torr, Hanrui Wang, Di Jin. Eigen-1: Adaptive Multi-Agent Refinement with Monitor-Based RAG for Scientific Reasoning. Preprint.

Wenxuan Huang, Shuang Chen, Zheyong Xie, Shaosheng Cao<sup>‡</sup>, Shixiang Tang, Yufan Shen, Qingyu Yin, Wenbo Hu, Xiaoman Wang, Yuntian Tang, Junbo Qiao, Yue Guo, Yao Hu, Zhenfei Yin<sup>‡</sup>, Philip Torr, Yu Cheng, Wanli Ouyang, Shaohui Lin<sup>‡</sup>. Interleaving Reasoning for Better Text-to-Image Generation. Preprint.

Zhuqiang Lu, Zhenfei Yin<sup>‡</sup>, Mengwei He, Zhihui Wang, Zicheng Liu, Zhiyong Wang, Kun Hu<sup>‡</sup>. B-VLLM: A Vision Large Language Model with Balanced Spatio-Temporal Tokens. International Conference on Computer Vision, ICCV 2025.

Xindi Yang, Baolu Li, Yiming Zhang, Zhenfei Yin<sup>‡</sup>, Lei Bai<sup>‡</sup>, Liqian Ma, Zhiyong Wang, Jianfei Cai, Tien-Tsin Wong, Huchuan Lu, Xu Jia<sup>‡</sup>. VLIPP: Towards Physically Plausible Video Generation with Vision and Language Informed Physical Prior. International Conference on Computer Vision, ICCV 2025.

Yiran Qin\*, Li Kang\*, Xiufeng Song\*, Zhenfei Yin<sup>‡</sup>, Xiaohong Liu, Xihui Liu, Ruimao Zhang<sup>‡</sup>, Lei Bai<sup>‡</sup>. RoboFactory: Exploring Embodied Agent Collaboration with Compositional Constraints. International Conference on Computer Vision, ICCV 2025.

Renqi Chen\*, Haoyang Su\*, Shixiang Tang, Zhenfei Yin, Qi Wu, Hui Li, Ye Sun, Nanqing Dong<sup>‡</sup>, Wanli Ouyang, Philip Torr. AI-Driven Automation Can Become the Foundation of Next-Era Science of Science Research. Preprint.

Rui Ye, Xiangrui Liu, Qimin Wu, Xianghe Pang, Zhenfei Yin, Lei Bai, Siheng Chen<sup>‡</sup>. X-MAS: Towards Building Multi-Agent Systems with Heterogeneous LLMs. Preprint.

Zonghao Ying, Yangguang Shao, Jianle Gan, Gan Xu, Junjie Shen, Wenxin Zhang, Quanchen Zou, Junzheng Shi, Zhenfei Yin, Mingchuan Zhang, Aishan Liu, Xianglong Liu. SecureWebArena: A Holistic Security Evaluation Benchmark for LVLM-based Web Agents. Preprint.

Qibing Ren, Sitao Xie, Longxuan Wei, [Zhenfei Yin](#), Junchi Yan, Lizhuang Ma, Jing Shao. When Autonomy Goes Rogue: Preparing for Risks of Multi-Agent Collusion in Social Systems. Preprint.

Sha Zhang, Suorong Yang, Tong Xie, Xiangyuan Xue, Zixuan Hu, Rui Li, Wenxi Qu, [Zhenfei Yin](#), Tianfan Fu, Di Hu, Andres M. Bran, Nian Ran, Bram Hoex, Wangmeng Zuo, Philippe Schwaller, Wanli Ouyang, Lei Bai, Yanyong Zhang, Lingyu Duan, Shixiang Tang<sup>‡</sup>, Dongzhan Zhou<sup>‡</sup>. Position: Intelligent Science Laboratory Requires the Integration of Cognitive and Embodied AI. Preprint.

2024

Rui Ye\*, Xianghe Pang\*, Jingyi Chai, Jiaao Chen, [Zhenfei Yin](#), Zhen Xiang, Xiaowen Dong, Jing Shao, Siheng Chen<sup>‡</sup>. Are We There Yet? Revealing the Risks of Utilizing Large Language Models in Scholarly Peer Review. Preprint. arXiv: [2412.01708](#) [cs.CL]<sup>↗</sup>

Ziyi Yang\*, Zaibin Zhang\*, Zirui Zheng, Yuxian Jiang, Ziyue Gan, Zhiyu Wang, Zijian Ling, Jinsong Chen, Martz Ma, Bowen Dong, Prateek Gupta, Shuyue Hu, [Zhenfei Yin](#)<sup>‡</sup>, Guohao Li<sup>‡</sup>, Xu Jia, Lijun Wang, Bernard Ghanem, Huchuan Lu, Wanli Ouyang, Yu Qiao, Philip Torr, Jing Shao<sup>‡</sup>. OASIS: Open Agents Social Interaction Simulations on One Million Agents. Neural Information Processing Systems Workshop on Open-World Agents, NeurIPS-W 2024, arXiv: [2411.11581](#) [cs.AI]<sup>↗</sup>

Yiran Qin\*, Zhelun Shi\*, Jiwen Yu, Xijun Wang, Enshen Zhou, Lijun Li, [Zhenfei Yin](#)<sup>†</sup>, Xihui Liu, Lu Sheng, Jing Shao<sup>‡</sup>, Lei Bai<sup>‡</sup>, Wanli Ouyang, Ruimao Zhang<sup>‡</sup>. WorldSimBench: Towards Video Generation Models as World Simulators. Technical Report. arXiv: [2410.18072](#) [cs.CV]<sup>↗</sup>

Haoyang Su\*, Renqi Chen\*, Shixiang Tang<sup>‡</sup>, Xinzhe Zheng, Jingzhe Li, [Zhenfei Yin](#), Wanli Ouyang, Nanqing Dong<sup>‡</sup>. Two Heads Are Better Than One: A Multi-Agent System Has the Potential to Improve Scientific Idea Generation. Preprint. arXiv: [2410.09403](#) [cs.AI]<sup>↗</sup>

Yisong Xiao, Aishan Liu, QianJia Cheng, [Zhenfei Yin](#), Siyuan Liang, Jiapeng Li, Jing Shao, Xianglong Liu, Dacheng Tao. GenderBias-VL: Benchmarking Gender Bias in Vision Language Models via Counterfactual Probing. Preprint. arXiv: [2407.00600](#) [cs.CV]<sup>↗</sup>

Yongting Zhang\*, Lu Chen\*, Guodong Zheng, Yifeng Gao, Rui Zheng, Jinlan Fu, [Zhenfei Yin](#), Senjie Jin, Yu Qiao, Xuanjing Huang, Feng Zhao, Tao Gui<sup>‡</sup>, Jing Shao<sup>‡</sup>. SPA-VL: A Comprehensive Safety Preference Alignment Dataset for Vision Language Model. Preprint. arXiv: [2406.12030](#) [cs.CV]<sup>↗</sup>

Zeren Chen\*, Zhelun Shi\*, Xiaoya Lu\*, Lehan He\*, Sucheng Qian, [Zhenfei Yin](#)<sup>†</sup>, Wanli Ouyang, Jing Shao<sup>‡</sup>, Yu Qiao, Cewu Lu, Lu Sheng<sup>‡</sup>. RH2oT-P: A Primitive-Level Robotic Manipulation Dataset Towards Composable Generalization Agents in Real-world Scenarios. Neural Information Processing Systems Workshop on Open-World Agents, NeurIPS-W 2024, arXiv: [2403.19622](#) [cs.RO]<sup>↗</sup>

Zhelun Shi\*, Zhipin Wang\*, Hongxing Fan\*, Zaibin Zhang, Lijun Li, Yongting Zhang, [Zhenfei Yin](#), Lu Sheng<sup>‡</sup>, Yu Qiao, Jing Shao<sup>‡</sup>. Assessment of Multimodal Large Language Models in Alignment with Human Values. Preprint. arXiv: [2403.17830](#) [cs.CV]<sup>↗</sup>

Enshen Zhou\*, Yiran Qin\*, [Zhenfei Yin](#), Yuzhou Huang, Ruimao Zhang<sup>‡</sup>, Lu Sheng<sup>‡</sup>, Yu Qiao, Jing Shao<sup>†</sup>. Chain-of-Imagination for Reliable Instruction Following in Decision Making. Neural Information Processing Systems Workshop on Open-World Agents, NeurIPS-W 2024, arXiv: [2403.12037](#) [cs.CV]<sup>↗</sup>

Chen Qian\*, Jie Zhang\*, Wei Yao\*, Dongrui Liu, [Zhenfei Yin](#), Yu Qiao, Yong Liu<sup>‡</sup>, Jing Shao<sup>‡</sup>. Towards Tracing Trustworthiness Dynamics: Revisiting Pre-training Period of Large Language Models. Findings of the Association for Computational Linguistics, ACL 2024, pp. 4864-4888. doi: [10.18653/v1/2024.findings-acl.290](#)<sup>↗</sup>. arXiv: [2402.19465](#) [cs.CL]<sup>↗</sup>

Chaochao Lu, Chen Qian, Guodong Zheng, Hongxing Fan, Hongzhi Gao, Jie Zhang, Jing Shao<sup>‡</sup>, Jingyi Deng, Jinlan Fu, Kexin Huang, Kunchang Li, Lijun Li, Limin Wang, Lu Sheng, Meiqi Chen, Ming Zhang, Qibing Ren, Sirui Chen, Tao Gui, Wanli Ouyang, Yali Wang, Yan Teng, Yaru Wang, Yi Wang, Yinan He, Yingchun Wang, Yixu Wang, Yongting Zhang, Yu Qiao<sup>‡</sup>, Yujiong Shen, Yurong Mou, Yuxi Chen, Zaibin Zhang, Zhelun Shi, Zhenfei Yin<sup>†</sup>, Zhipin Wang. From GPT-4 to Gemini and Beyond: Assessing the Landscape of MLLMs on Generalizability, Trustworthiness and Causality through Four Modalities. Technical Report. arXiv: 2401.15071 [cs.CV]<sup>↗</sup>

2023

Zhiyuan You\*, Zheyuan Li\*, Jinjin Gu\*, Zhenfei Yin, Tianfan Xue<sup>‡</sup>, Chao Dong<sup>‡</sup>. Depicting beyond scores: Advancing image quality assessment through multi-modal language models. European Conference on Computer Vision, ECCV 2024, pp. 259-276. doi: 10.1007/978-3-031-72970-6\_15<sup>↗</sup>. arXiv: 2312.08962 [cs.CV]<sup>↗</sup>

Yiran Qin\*, Enshen Zhou\*, Qichang Liu\*, Zhenfei Yin, Lu Sheng<sup>‡</sup>, Ruimao Zhang<sup>‡</sup>, Yu Qiao, Jing Shao<sup>†</sup>. MP5: A Multi-Modal Open-Ended Embodied System in Minecraft via Active Perception. IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, pp. 16307-16316. doi: 10.1109/CVPR52733.2024.01543<sup>↗</sup>. arXiv: 2312.07472 [cs.CV]<sup>↗</sup>

Zhelun Shi\*, Zhipin Wang\*, Hongxing Fan\*, Zhenfei Yin, Lu Sheng<sup>‡</sup>, Yu Qiao, Jing Shao<sup>‡</sup>. ChEF: A Comprehensive Evaluation Framework for Standardized Assessment of Multimodal Large Language Models. Technical Report. arXiv: 2311.02692 [cs.CV]<sup>↗</sup>

Zeren Chen\*, Ziqin Wang\*, Zhen Wang, Huayang Liu, Zhenfei Yin<sup>†</sup>, Si Liu, Lu Sheng<sup>‡</sup>, Wanli Ouyang, Jing Shao<sup>‡</sup>. Octavius: Mitigating Task Interference in MLLMs via LoRA-MoE. International Conference on Learning Representations, ICLR 2024, arXiv: 2311.02684 [cs.CV]<sup>↗</sup>

Zhenfei Yin\*, Jiong Wang\*, Jianjian Cao\*, Zhelun Shi\*, Dingning Liu, Mukai Li, Xiaoshui Huang, Zhiyong Wang, Lu Sheng, Lei Bai<sup>‡</sup>, Jing Shao<sup>‡</sup>, Wanli Ouyang. LAMM: Language-Assisted Multi-Modal Instruction-Tuning Dataset, Framework, and Benchmark. Advances in Neural Information Processing Systems, NeurIPS 2023, Vol. 36. arXiv: 2306.06687 [cs.CV]<sup>↗</sup>

2022

Yuan Yao, Yuanhan Zhang, Zhenfei Yin, Jiebo Luo, Wanli Ouyang, Xiaoshui Huang<sup>‡</sup>. 3D Point Cloud Pre-Training with Knowledge Distilled from 2D Images. IEEE International Conference on Multimedia and Expo, ICME 2024, pp. 1-6. doi: 10.1109/ICME57554.2024.10687626<sup>↗</sup>. arXiv: 2212.08974 [cs.CV]<sup>↗</sup>

Yuanhan Zhang, Zhenfei Yin<sup>†</sup>, Jing Shao<sup>‡</sup>, Ziwei Liu. Benchmarking omni-vision representation through the lens of visual realms. European Conference on Computer Vision, ECCV 2022, pp. 594-611. doi: 10.1007/978-3-031-20071-7\_35<sup>↗</sup>. arXiv: 2207.07106 [cs.CV]<sup>↗</sup>

Yinan He\*, Gengshi Huang\*, Siyu Chen\*, Jianing Teng\*, Kun Wang, Zhenfei Yin, Lu Sheng, Ziwei Liu, Yu Qiao, Jing Shao<sup>‡</sup>. X-learner: Learning cross sources and tasks for universal visual representation. European Conference on Computer Vision, ECCV 2022, pp. 509-528. doi: 10.1007/978-3-031-19809-0\_29<sup>↗</sup>. arXiv: 2203.08764 [cs.CV]<sup>↗</sup>

Yuanhan Zhang\*, Qinghong Sun\*, Yichun Zhou\*, Zexin He\*, Zhenfei Yin<sup>†</sup>, Kun Wang, Lu Sheng, Yu Qiao, Jing Shao<sup>‡</sup>, Ziwei Liu. Bamboo: Building Mega-Scale Vision Dataset Continually with Human-Machine Synergy. International Journal of Computer Vision, 2025: 1-16

2021

Yujie Wang\*, Junqin Huang\*, Mengya Gao\*, Yichao Wu\*, Zhenfei Yin, Ding Liang, Junjie Yan. One to Transfer All: A Universal Transfer Framework for Vision Foundation Model with Few Data. Technical Report. arXiv: 2111.12386 [cs.CV]<sup>↗</sup>

Jing Shao\*, Siyu Chen\*, Yangguang Li\*, Kun Wang\*, [Zhenfei Yin\\*](#), Yanan He\*, Jianing Teng\*, Qinghong Sun\*, Mengya Gao\*, Jihao Liu\*, Gengshi Huang\*, Guanglu Song, Yichao Wu, Yuming Huang, Fenggang Liu, Huan Peng, Shuo Qin, Chengyu Wang, Yujie Wang, Conghui He, Ding Liang, Yu Liu, Fengwei Yu, Junjie Yan, Dahua Lin, Xiaogang Wang, Yu Qiao<sup>‡</sup>. INTERN: A New Learning Paradigm Towards General Vision. Technical Report. arXiv: [2111.08687](#) [cs.CV]<sup>↗</sup>

## HONOURS

- 2025 WAIC 2025 Yunfan Award Rising Star nominee
- 2022 Rising Star at the Shanghai AI Lab
- 2020 4th Place Winner in Meta's DFDC Challenge
- 2018 Student coach for the team that reached the ACM-ICPC World Finals, finishing 31st globally

## INVITED TALKS & TEACHING

- 2024 Building AI Society with Foundation-Model Agents.  
Workshop on [Open-World Agents](#)<sup>↗</sup> at NeurIPS 2024  
Multi-Modal Foundation Models and Applications as Embodied Agents  
Guest Lecture, [ELEC5304](#)<sup>↗</sup>: Intelligent Visual Signal Understanding, USYD
- 2024 Spring Multimodal Applications  
Teaching Assistant, [COMP 5425](#)<sup>↗</sup>: Multimedia Retrieval, USYD

## WORKSHOP & ACADEMIC EVENT ORGANIZATION

- 2026 Workshop Organizer, \*Scalable Robot Learning Systems\* (FIRST Workshop), CVPR 2026.  
Workshop Organizer, \*Lifelong Agents: Learning, Aligning, Evolving\*, ICLR 2026.  
Workshop Organizer, \*AAVM: Agentic AI for Visual Media\*, CVPR 2026.  
Workshop Organizer, \*The 6th Workshop of Adversarial Machine Learning on Computer Vision: Safety of Vision-Language Agents\*, CVPR 2026.
- 2025 Competition Organizer, \*Multi-Agent Embodied Intelligence Challenge\*, ICML 2025 and NeurIPS 2025.  
Competition Organizer, \*SFE Challenge\* (AI Scientist Challenge), ICCV 2025.  
Workshop Organizer, \*Reliable and Interactable World Models: Geometry, Physics, Interactivity and Real-World Generalization\*, ICCV 2025.  
Workshop Organizer, \*Multi-Modal Reasoning for Agentic Intelligence\*, ICCV 2025.  
Workshop Organizer, \*Multi-Agent Systems in the Era of Foundation Models: Opportunities, Challenges, and Futures (MAS-2025)\*, ICML 2025.

- 2024 Workshop Organizer, ICML 2024 workshop on Multi-modal Foundation Model meets Embodied AI, [MFM-EAI](#)
- Workshop Organizer, ICML 2024 workshop on Trustworthy Multi-modal Foundation Models and AI Agents, [TiFA](#)
- Since 2023 Academic Talks Organizer, [Echo AI Talk](#)

## PROGRAMME COMMITTEE MEMBERSHIPS

### Conferences

- |   |            |
|---|------------|
| International Conference on Learning Representations (ICLR)           | Since 2023 |
| Neural Information Processing Systems (NeurIPS)                       | Since 2023 |
| International Conference on Machine Learning (ICML)                   | Since 2024 |
| Association for the Advancement of Artificial Intelligence (AAAI)     | Since 2023 |
| European Conference on Computer Vision (ECCV)                         | Since 2022 |
| IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) | Since 2022 |

### Journals

- |  |            |
|--|------------|
| IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) | Since 2023 |
|--|------------|