

王艺炜

电子邮箱: wangyw.evan@gmail.com 微信号: wangywseu 个人主页: <https://wangywust.github.io/>

工作经历

- | | |
|-----------------------------|-----------------|
| • 东南大学, 教授
信息科学与工程学院 | 2026 年 - 至今 |
| • 加州大学默塞德分校, 助理教授
计算机科学系 | 2025 年 - 2026 年 |
| • 加州大学洛杉矶分校, 博士后研究员 | 2023 年 - 2024 年 |
| • 亚马逊公司 (西雅图), 应用科学家 | 2022 年 - 2023 年 |

教育背景

- | | |
|---|-----------------|
| • 新加坡国立大学, 博士 | 2019 年 - 2023 年 |
| • 香港科技大学, 研究型硕士 (Master of Philosophy, MPhil) | 2017 年 - 2019 年 |
| • 东南大学信息科学与工程学院, 本科 | 2013 年 - 2017 年 |

研究论文发表

- [1] Baolong Bi, Yuyao Ge, Shenghua Liu, Yuchen He, Siqian Tong, Lizhe Chen, Lingrui Mei, Zehao Li, Yiwei Wang, Yujun Cai, Ming-Hsuan Yang, Xueqi Cheng. PromptCD: Test-Time Behavior Enhancement via Polarity-Prompt Contrastive Decoding, *TPAMI 2026*. (CCF-A) (JCR Q1)
- [2] Honghao Fu, Jiacheng Chen, Manxi Lin, Junjun Zheng, Xiangheng Kong, Yiwei Wang, Xin Yu, Miao Xu, Yuning Jiang, Yujun Cai. AdaM-Rec: Adaptive Modality Routing for Multimodal Recommendation, *NeurIPS 2026*. (CCF-A)
- [3] Yuyao Ge, Yiwei Wang, Yuchen He, Baolong Bi, Lingrui Mei, Jiayu Yao, Lizhe Chen, Shenghua Liu. SkillForge: Co-Evolving Skills and Agents via Dynamic Skill Lifecycles, *NeurIPS 2026*. (CCF-A)
- [4] Haonan Ge, Yiwei Wang, Hang Wu, Yujun Cai. What Should a Streaming Video Model Remember?, *NeurIPS 2026*. (CCF-A)
- [5] Zhixiong Nan, Haoyu Lu, Tao Xiang, Yiwei Wang. SCTI: Self-Calibrated Trident Identification of Black-Box LLM Watermarks, *NeurIPS 2026*. (CCF-A)
- [6] Rongzhi Zhu, Yi Liu, Jiancheng Wang, Xiangyu Liu, Zequn Sun, Yiwei Wang, Yu Deng, Zijian Zhou, Wei Hu. When Can Large Reasoning Models Save Thinking? Mechanistic Analysis of Behavioral Divergence in Reasoning, *EMNLP 2026*.
- [7] Jiayu Yao, Yiwei Wang, Anmeng Zhang, Zhe Sun, Songsong Wang, Lingrui Mei, Yuyao Ge, Shenghua Liu. Multimodal Reward Hacking in Reinforcement Learning, *EMNLP 2026*.
- [8] Lexin Wang, Shenghua Liu, Yiwei Wang, Jiafeng Guo, Xueqi Cheng. Readable Yet Unpredictable: Rotated-Outcome Prediction in Vision-Language Models, *EMNLP 2026*.
- [9] Zehao Li, Hao Jiang, Yujun Cai, Jianing Chen, Baolong Bi, Shuqin Gao, Honglong Zhao, Yiwei Wang, Tianlu Mao, Zhaoqi Wang. STDR: Spatio-Temporal Decoupling for Real-Time Dynamic Scene Rendering, *ACM MM 2026*. (CCF-A)
- [10] Zhetong Zhang, Honghao Fu, Miao Xu, Yiwei Wang, Yujun Cai. V-DEAL: Diagnosing Video Safety De-Calibration as an Understanding-Refusal Coupling Failure, *COLM 2026*.

- [11] Yuyao Ge, Shenghua Liu, Yiwei Wang, Baolong Bi, Lingrui Mei, Jiayu Yao, Tianyu Liu, Jiafeng Guo, Xueqi Cheng. Prism-: Differential Subspace Steering for Prompt Highlighting in Large Language Models, *COLM 2026*.
- [12] Lexin Wang, Shenghua Liu, Yiwei Wang, Yujun Cai, Yuyao Ge, Jiayu Yao, Jiafeng Guo, Xueqi Cheng. HighlightBench: Benchmarking and Diagnosing Markup-Driven Table Reasoning in Scientific Documents, *ECCV 2026*.
- [13] Yuyao Ge, Shenghua Liu, Yiwei Wang, Lingrui Mei, Baolong Bi, Xuanshan Zhou, Jiayu Yao, Jiafeng Guo, Xueqi Cheng. Focusing by Contrastive Attention: Enhancing VLMs’ Visual Reasoning, *ECCV 2026*.
- [14] Xingjian Tao, Yiwei Wang, Yujun Cai, Yifan Song, Jing Tang. ViewFusion: Structured Spatial Thinking Chains for Multi-View Reasoning, *ECCV 2026*.
- [15] Xingjian Tao, Yiwei Wang, Yujun Cai, Yihong Luo, Kai Han, Jing Tang. Mitigating Coordinate Prediction Bias from Positional Encoding Failures, *ACL 2026*. (CCF-A)
- [16] Boyan Han, Yiwei Wang, Yi Song, Yujun Cai, Chi Zhang. Dynamic Infilling Anchors for Format-Constrained Generation in Diffusion Large Language Models, *ACL 2026*. (CCF-A)
- [17] Zhixue Song, Boyan Han, Yiwei Wang, Chi Zhang. Hard to Read, Easy to Jailbreak: How Visual Degradation Bypasses MLLM Safety Alignment, *ACL 2026*. (CCF-A)
- [18] Xingjian Tao, Yiwei Wang, Yujun Cai, Yihong Luo, Jing Tang. Rectifying Coordinate Drift in MLLMs via Counterfactual Positional Guidance, *ACL 2026*. (CCF-A)
- [19] Zheng He, Yiwei Wang, Hongxing Wang, Yujun Cai. ReCon: Active Defense against Large Vision-Language Model Jailbreaks via Reverse Safety Concept Injection, *ACL 2026*. (CCF-A)
- [20] Lingrui Mei, Shenghua Liu, Yiwei Wang, Yuyao Ge, Baolong Bi, Jiayu Yao, Jun Wan, Ziling Yin, Jiafeng Guo, Xueqi Cheng. Gated Differentiable Working Memory for Long-Context Language Modeling, *ACL 2026*. (CCF-A)
- [21] Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Ruibin Yuan, Xueqi Cheng. HiddenGuard: Fine-Grained Safe Generation with Specialized Representation Router, *ACL 2026*. (CCF-A)
- [22] Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Yuyao Ge, Jun Wan, Yurong Wu, Xueqi Cheng. a1: Steep Test-time Scaling Law via Environment Augmented Generation, *ACL 2026*. (CCF-A)
- [23] Honghao Fu, Miao Xu, Yiwei Wang, Dailing Zhang, Jun Liu, Yujun Cai. VideoStir: Understanding Long Videos via Spatio-Temporally Structured and Intent-Aware RAG, *ACL 2026*. (CCF-A)
- [24] Wenjia Jiang, Yiwei Wang, Boyan Han, Joey Tianyi Zhou, Chi Zhang. SQLAgent: Learning to Explore Before Generating as a Data Engineer, *ACL 2026*. (CCF-A)
- [25] Zhicheng Yang, Zhijiang Guo, Yinya Huang, Yongxin Wang, Dongchun Xie, Hanhui Li, Yiwei Wang, Xiaodan Liang, Jing Tang. Depth-Breadth Synergy in RLVR: Unlocking LLM Reasoning Gains with Adaptive Exploration, *ICML 2026*. (CCF-A)
- [26] Siqian Tong, Li Xuan, Yiwei Wang, Baolong Bi, Yujun Cai, Shenghua Liu, Yuchen He, HAO Chengpeng. AuTAgent: A Reinforcement Learning Framework for Tool-Augmented Audio Reasoning, *ICML 2026*. (CCF-A)
- [27] Baolong Bi, Shenghua Liu, Yiwei Wang, Siqian Tong, Lingrui Mei, Yuyao Ge, Yilong Xu, Jiafeng Guo, Xueqi Cheng. Reward and Guidance through Rubrics: Promoting Exploration to Improve Multi-Domain Reasoning, *ICML 2026*. (CCF-A)
- [28] Yixin Song, Guangchi Liu, Xiangyu Xu, Shaofeng Li, Zhen Ling, Yiwei Wang, Yujun Cai. MIMO-LP: A Multi-Input Multi-Output Framework for Subgraph-based Link Prediction, *ICML 2026*. (CCF-A)
- [29] Zhicheng Yang, Zhijiang Guo, Yinya Huang, Yongxin Wang, Wenlei Shi, Yiwei Wang, Xiaodan Liang, Jing Tang. Accordion-Thinking: Self-Regulated Step Summaries for Efficient and Readable LLM Reasoning, *ICML 2026*. (CCF-A)

- [30] Bowen Sun, Yujun Cai, Ming-Hsuan Yang, Hang Wu, Yiwei Wang. PAS: A Training-Free Stabilizer for Temporal Encoding in Video LLMs, *CVPR 2026*. (CCF-A)
- [31] Samyak Rawlekar, Amitabh Swain, Yujun Cai, Yiwei Wang, Ming-Hsuan Yang, Narendra Ahuja. Finding Distributed Object-Centric Properties in Self-Supervised Transformers, *CVPR 2026*. (CCF-A)
- [32] Liangyu Yuan, Yufei splinter Huang, Mingkun Lei, Tong Zhao, Ruoyu Wang, Changxi Chi, Yiwei Wang, Chi Zhang. Improving Diffusion Generalization with Weak-to-Strong Segmented Guidance, *CVPR 2026*. (CCF-A)
- [33] Bingxuan Li, Yiming Cui, Yicheng He, Yiwei Wang, Shu Zhang, Longyin Wen, Yulei Niu. EchoFoley: Event-Centric Hierarchical Control for Video Grounded Creative Sound Generation, *CVPR 2026*. (CCF-A)
- [34] Zhenlong Yuan, Xiangyan Qu, Chengxuan Qian, Rui Chen, Jing Tang, Lei Sun, Xiangxiang Chu, Dapeng Zhang, Yiwei Wang, Yujun Cai, Shuo Li. Video-STAR: Reinforcing Open-Vocabulary Action Recognition with Tools, *ICLR 2026*. (CCF-A)
- [35] Honghao Fu, Yuan Ouyang, Kai-Wei Chang, Yiwei Wang, Zi Huang, Yujun Cai. ContextNav: Towards Agentic Multimodal In-Context Learning, *ICLR 2026*. (CCF-A)
- [36] Haojin Yang, Rui Hu, Zequn Sun, Rui Zhou, Yujun Cai, Yiwei Wang. WavefrontDiffusion: Dynamic Decoding Schedule for Improved Reasoning, *ICLR 2026*. (CCF-A)
- [37] Baolong Bi, Shenghua Liu, Yiwei Wang, Yilong Xu, Junfeng Fang, Lingrui Mei, Xueqi Cheng. Parameters vs. Context: Fine-Grained Control of Knowledge Reliance in Language Models, *ICLR 2026*. (CCF-A)
- [38] Yike Wu, Yiwei Wang, Yujun Cai. ChainMPQ: Interleaved Text-Image Reasoning Chains for Mitigating Relation Hallucinations, *ICLR 2026*. (CCF-A)
- [39] Xingjian Tao, Yiwei Wang, Yujun Cai, Zhicheng Yang, Jing Tang. Are LLMs Really Not Knowledgeable? Mining the Submerged Knowledge in LLMs’ Memory, *ICLR 2026*. (CCF-A)
- [40] Xuan Shen, Weize Ma, Yufa Zhou, Enhao Tang, Yanyue Xie, Zhengang Li, Yifan Gong, Quanyi Wang, Henghui Ding, Yiwei Wang, Pu Zhao, Jun Lin, Jiuxiang Gu. Fastcar: Cache Attentive Replay for Fast Auto-Regressive Video Generation on the Edge, *ICLR 2026*. (CCF-A)
- [41] Zhen Xiong, Yujun Cai, Zhecheng Li, Yiwei Wang. Unveiling the Potential of Diffusion Large Language Model in Controllable Generation, *ICLR 2026*. (CCF-A)
- [42] Xuan Shen, Brian Wingenroth, Zichao Wang, Jason Kuen, Wanrong Zhu, Ruiyi Zhang, Yiwei Wang, Lichun Ma, Anqi Liu, Hongfu Liu, Tong Sun, Kevin S. Hawkins, Kate Tasker, G. Caleb Alexander, Jiuxiang Gu. OIDA-QA: A Multimodal Benchmark for Analyzing the Opioid Industry Documents Archive, *AAAI 2026*. (CCF-A)
- [43] Haodong Xie, Yujun Cai, Rahul Singh Maharjan, Yiwei Wang, Federico Tavella, Angelo Cangelosi. Hierarchical, Interpretable, Label-Free Concept Bottleneck Model, *British Machine Vision Conference (BMVC) 2026*.
- [44] Jiayu Yao, Shenghua Liu, Yiwei Wang, Rundong Cheng, Lingrui Mei, Baolong Bi, Zhen Xiong, Xueqi Cheng. Not in Sync: Unveiling Temporal Bias in Large Audio-Language Models, *China Conference on Artificial Intelligence (CICAI) 2026*.
- [45] Hui Chen, James Xu Zhao, Dongfu Jiang, Qianyun Guo, Jiefeng Chen, Yiwei Wang, Muhao Chen, See-Kiong Ng, Pang Wei Koh, Bryan Hooi. FabScore: Fine-Grained Evaluation of Fabrications in Automated AI Research, *ICML 2026 Workshop on AI for Science 2026*.
- [46] Yao Gao, Junlei Hu, Feng Li, Xijin Du, Yiwei Wang, Jeroen Van Dessel, Robin Willaert, Yi Sun. OrbitPSIBench: Benchmarking Vision-Free LLMs’ Anatomical Reasoning in Orbital Bone, *Medical Image Computing and Computer Assisted Intervention (MICCAI) 2026*.
- [47] Yifan Lu, Pengfei Jiao, Xuan Guo, Ziyun Zou, Yiwei Wang, Mengzhou Gao, Huaming Wu, Imran Razzak.

- Towards Robust Heterogeneous Graph Explanations under Structural Perturbations, *WWW 2026*. (CCF-A)
- [48] Juncheng Liu, Chenghao Liu, Gerald Woo, Yiwei Wang, Bryan Hooi, Caiming Xiong, Doyen Sahoo. UniTST: Effectively Modeling Inter-Series and Intra-Series Dependencies for Multivariate Time Series Forecasting, *TMLR 2025*.
- [49] Jiayu Yao, Shenghua Liu, Yiwei Wang, Lingrui Mei, Baolong Bi, Yuyao Ge, Zhecheng Li, Xueqi Cheng. Who is in the Spotlight: The Hidden Bias Undermining Multimodal Retrieval-Augmented Generation, *EMNLP 2025*.
- [50] Zhen Xiong, Yujun Cai, Zhecheng Li, Yiwei Wang. Mapping the Minds of LLMs: A Graph-Based Analysis of Reasoning LLMs, *EMNLP 2025*.
- [51] Shuyang Hao, Yiwei Wang, Bryan Hooi, Jun Liu, Muhao Chen, Zi Huang, Yujun Cai. Making every step effective: Jailbreaking large vision-language models through hierarchical kv equalization, *EMNLP 2025*.
- [52] Haonan Ge, Yiwei Wang, Ming-Hsuan Yang, Yujun Cai. MRFD: Multi-Region Fusion Decoding with Self-Consistency for Mitigating Hallucinations in LVLMs, *EMNLP 2025*.
- [53] Hang Wu, Hongkai Chen, Yujun Cai, Chang Liu, Qingwen Ye, Ming-Hsuan Yang, Yiwei Wang. DiMo-GUI: Advancing Test-time Scaling in GUI Grounding via Modality-Aware Visual Reasoning, *EMNLP 2025*.
- [54] Wen Tao, Yiwei Wang. How to Make Large Language Models Generate 100% Valid Molecules?, *EMNLP 2025*.
- [55] Sifan Li, Yujun Cai, Yiwei Wang. SemVink: Advancing VLMs' Semantic Understanding of Optical Illusions via Visual Global Thinking, *EMNLP 2025*.
- [56] Xingjian Tao, Yiwei Wang, Yujun Cai, Zhicheng Yang, Jing Tang. Understanding GUI Agent Localization Biases through Logit Sharpness, *EMNLP 2025*.
- [57] Zhecheng Li, Guoxian Song, Yujun Cai, Zhen Xiong, Junsong Yuan, Yiwei Wang. Texture or Semantics? Vision-Language Models Get Lost in Font Recognition, *COLM 2025*.
- [58] Zhaochen Wang, Bryan Hooi, Yiwei Wang, Ming-Hsuan Yang, Zi Huang, Yujun Cai. Text Speaks Louder than Vision: ASCII Art Reveals Textual Biases in Vision-Language Models, *COLM 2025*.
- [59] Chunxue Xu, Yiwei Wang, Bryan Hooi, Yujun Cai, Songze Li. How does Watermarking Affect Visual Language Models in Document Understanding?, *COLM 2025*.
- [60] Baolong Bi, Shaohan Huang, Yiwei Wang, Tianchi Yang, Zihan Zhang, Haizhen Huang, Lingrui Mei, Junfeng Fang, Zehao Li, Furu Wei, Weiwei Deng, Feng Sun, Qi Zhang, Shenghua Liu. Context-DPO: Aligning Language Models for Context-Faithfulness, *ACL 2025*. (CCF-A)
- [61] Baolong Bi, Shenghua Liu, Lingrui Mei, Yiwei Wang, Junfeng Fang, Pengliang Ji, Xueqi Cheng. Decoding by Contrasting Knowledge: Enhancing Large Language Model Confidence on Edited Facts, *ACL 2025*. (CCF-A)
- [62] Yuyao Ge, Shenghua Liu, Baolong Bi, Yiwei Wang, Lingrui Mei, Wenjie Feng, Lizhe Chen, Xueqi Cheng. Can Graph Descriptive Order Affect Solving Graph Problems with LLMs?, *ACL 2025*. (CCF-A)
- [63] Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Nanyun Peng, Kai-Wei Chang. DRS: Deep Question Reformulation With Structured Output, *ACL 2025*. (CCF-A)
- [64] Bingxuan Li, Yiwei Wang, Jiuxiang Gu, Kai-Wei Chang, Nanyun Peng. METAL: A Multi-Agent Framework for Chart Generation with Test-Time Scaling, *ACL 2025*. (CCF-A)
- [65] Zhecheng Li, Yiwei Wang, Bryan Hooi, Yujun Cai, Zhen Xiong, Nanyun Peng, Kai-Wei Chang. Vulnerability of LLMs to Vertically Aligned Text Manipulations, *ACL 2025*. (CCF-A)
- [66] Yurong Wu, Fangwen Mu, Qihong Zhang, Jinjing Zhao, Xinrun Xu, Lingrui Mei, Yang Wu, Lin Shi, Junjie Wang, Zhiming Ding, Yiwei Wang. EvoStealer: Differential Evolution for Prompt Template Stealing Against Text-to-Image Synthesis, *ACL 2025*. (CCF-A)

- [67] Rongzhi Zhu, Xiangyu Liu, Zequn Sun, Yiwei Wang, Wei Hu. Mitigating Lost-in-Retrieval Problems in Retrieval Augmented Multi-Hop Question Answering, *ACL 2025*. (CCF-A)
- [68] Songlin Zhai, Yuan Meng, Yongrui Chen, Yiwei Wang, Guilin Qi. Peripheral Memory for LLMs: Integration of Sequential Memory Banks with Adaptive Querying, *ICML 2025*. (CCF-A)
- [69] Baolong Bi, Shenghua Liu, Yiwei Wang, Lingrui Mei, Xueqi Cheng. Is Factuality Enhancement a Free Lunch For LLMs? Better Factuality Can Lead to Worse Context-Faithfulness, *ICLR 2025*. (CCF-A)
- [70] Zhicheng Yang, Yiwei Wang, Yinya Huang, Zhijiang Guo, Wei Shi, Xiongwei Han, Liang Feng, Linqi Song, Xiaodan Liang, Jing Tang. OptiBench meets ReSocratic: Measure and improve LLMs for optimization modeling, *ICLR 2025*. (CCF-A)
- [71] Shaochen Zhong, others. MQuAKE-Remastered: Multi-Hop Knowledge Editing Can Only Be Advanced With Reliable Evaluations, *ICLR 2025*. (CCF-A)
- [72] Cheng Wang, Yiwei Wang, Yujun Cai, Bryan Hooi. Tricking Retrievers with Influential Tokens: An Efficient Black-Box Corpus Poisoning Attack, *NAACL 2025*.
- [73] Yiwei Wang, Muhao Chen, Nanyun Peng, Kai-Wei Chang. Vulnerability of Large Language Models to Output Prefix Jailbreaks: Impact of Positions on Safety, *NAACL 2025*.
- [74] Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Jiayi Mao, Xueqi Cheng. "Not Aligned" is Not "Malicious": Being Careful about Hallucinations of Large Language Models' Jailbreak, *COLING 2025*.
- [75] Cheng Wang, Yiwei Wang, Bryan Hooi, Yujun Cai, Nanyun Peng, Kai-Wei Chang. Con-ReCall: Detecting Pre-training Data in LLMs via Contrastive Decoding, *COLING 2025*.
- [76] Yihong Luo, Yuhao Chen, Siya Qiu, Yiwei Wang, Chen Zhang, Yan Zhou, Xiaochun Cao, Jing Tang. Fast Graph Sharpness-Aware Minimization for Enhancing and Accelerating Few-Shot Node Classification, *NeurIPS 2024*. (CCF-A)
- [77] Bingxuan Li, Yiwei Wang, Tao Meng, Nanyun Peng, Kai-Wei Chang. Control Large Language Models via Divide-and-Conquer, *EMNLP 2024*.
- [78] Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Xueqi Chen. SLANG: New Concept Comprehension of Large Language Models, *EMNLP 2024*.
- [79] Silin Meng, Yiwei Wang, Cheng-Fu Yang, Nanyun Peng, Kai-Wei Chang. LLM-A*: Large Language Model Enhanced Incremental Heuristic Search on Path Planning, *EMNLP 2024*.
- [80] Zhicheng Yang, Yinya Huang, Jing Xiong, Liang Feng, Xiaodan Liang, Yiwei Wang, Jing Tang. AlignedCoT: Prompting Large Language Models via Native-Speaking Demonstrations, *EMNLP 2024*.
- [81] Baolong Bi, Shenghua Liu, Yiwei Wang, Lingrui Mei, Xueqi Cheng. Lpnl: Scalable link prediction with large language models, *ACL 2024*. (CCF-A)
- [82] Juncheng Liu, Bryan Hooi, Kenji Kawaguchi, Yiwei Wang, Chaosheng Dong, Xiaokui Xiao. Scalable and Effective Implicit Graph Neural Networks on Large Graphs, *ICLR 2024*. (CCF-A)
- [83] Fei Wang, Wenjie Mo, Yiwei Wang, Wenxuan Zhou, Muhao Chen. A Causal View of Entity Bias in (Large) Language Models, *EMNLP 2023*.
- [84] Yiwei Wang, Yujun Cai, Muhao Chen, Yuxuan Liang, Bryan Hooi. Primacy Effect of ChatGPT, *EMNLP 2023*.
- [85] Yuxuan Liang, Yutong Xia, Songyu Ke, Yiwei Wang, Qingsong Wen, Junbo Zhang, Yu Zheng, Roger Zimmermann. Airformer: Predicting nationwide air quality in china with transformers, *AAAI 2023*. (CCF-A)
- [86] Yiwei Wang, Bryan Hooi, Fei Wang, Yujun Cai, Yuxuan Liang, Wenxuan Zhou, Jing Tang, Manjuan Duan, Muhao Chen. How Fragile is Relation Extraction under Entity Replacements?, *CONLL 2023*.

- [87] Yiwei Wang, Bryan Hooi, Yozen Liu, Neil Shah. Graph explicit neural networks: Explicitly encoding graphs for efficient and accurate inference, *WSDM 2023*.
- [88] Yuxuan Liang, Kun Ouyang, Yiwei Wang, Zheyi Pan, Yifang Yin, Hongyang Chen, Junbo Zhang, Yu Zheng, David S Rosenblum, Roger Zimmermann. Mixed-order relation-aware recurrent neural networks for spatio-temporal forecasting, *TKDE 2022*. (CCF-A) (JCR Q1)
- [89] Juncheng Liu, Zequn Sun, Bryan Hooi, Yiwei Wang, Dayiheng Liu, Baosong Yang, Xiaokui Xiao, Muhao Chen. Dangling-Aware Entity Alignment with Mixed High-Order Proximities, *NAACL 2022*.
- [90] Yiwei Wang, Muhao Chen, Wenxuan Zhou, Yujun Cai, Yuxuan Liang, Bryan Hooi. GraphCache: Message Passing as Caching for Sentence-Level Relation Extraction, *NAACL 2022*.
- [91] Yiwei Wang, Muhao Chen, Wenxuan Zhou, Yujun Cai, Yuxuan Liang, Dayiheng Liu, Baosong Yang, Juncheng Liu, Bryan Hooi. Should We Rely on Entity Mentions for Relation Extraction? Debiasing Relation Extraction with Counterfactual Analysis, *NAACL 2022*.
- [92] Juncheng Liu, Yiwei Wang, Bryan Hooi, Renchi Yang, Xiaokui Xiao. Lscale: latent space clustering-based active learning for node classification, *ECML 2022*.
- [93] Yiwei Wang, Bryan Hooi, Yozen Liu, Tong Zhao, Zhichun Guo, Neil Shah. Flashlight: Scalable link prediction with effective decoders, *LOG 2022*.
- [94] Juncheng Liu, Kenji Kawaguchi, Bryan Hooi, Yiwei Wang, Xiaokui Xiao. Eignn: Efficient infinite-depth graph neural networks, *NeurIPS 2021*. (CCF-A)
- [95] Yiwei Wang, Yujun Cai, Yuxuan Liang, Henghui Ding, Changhu Wang, Siddharth Bhatia, Bryan Hooi. Adaptive data augmentation on temporal graphs, *NeurIPS 2021*. (CCF-A)
- [96] Yujun Cai, Yiwei Wang, Yiheng Zhu, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Chuanxia Zheng, Sijie Yan, Henghui Ding, others. A unified 3d human motion synthesis model via conditional variational auto-encoder, *ICCV 2021*. (CCF-A)
- [97] Yiwei Wang, Yujun Cai, Yuxuan Liang, Henghui Ding, Changhu Wang, Bryan Hooi. Time-aware neighbor sampling for temporal graph networks, *2022 International Joint Conference on Neural Networks (IJCNN) 2021*.
- [98] Yuxuan Liang, Kun Ouyang, Yiwei Wang, Ye Liu, Junbo Zhang, Yu Zheng, David S Rosenblum. Revisiting convolutional neural networks for citywide crowd flow analytics, *ECML 2021*.
- [99] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, Bryan Hooi. Progressive supervision for node classification, *ECML 2021*.
- [100] Yuxuan Liang, Kun Ouyang, Hanshu Yan, Yiwei Wang, Zekun Tong, Roger Zimmermann. Modeling Trajectories with Neural Ordinary Differential Equations, *IJCAI 2021*. (CCF-A)
- [101] Yuxuan Liang, Kun Ouyang, Junkai Sun, Yiwei Wang, Junbo Zhang, Yu Zheng, David Rosenblum, Roger Zimmermann. Fine-grained urban flow prediction, *WWW 2021*. (CCF-A)
- [102] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, Bryan Hooi. Curgraph: Curriculum learning for graph classification, *WWW 2021*. (CCF-A)
- [103] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, Bryan Hooi. Mixup for node and graph classification, *WWW 2021*. (CCF-A)
- [104] Yujun Cai, Lin Huang, Yiwei Wang, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Xu Yang, Yiheng Zhu, Xiaohui Shen, others. Learning progressive joint propagation for human motion prediction, *ECCV 2020*.
- [105] Yuxuan Liang, Kun Ouyang, Yiwei Wang, David Samuel Rosenblum. Revisiting convolutional neural networks for urban flow analytics, *2020*.

- [106] Yiwei Wang, Shenghua Liu, Minji Yoon, Hemank Lamba, Wei Wang, Christos Faloutsos, Bryan Hooi. Provably robust node classification via low-pass message passing, *ICDM 2020*.
- [107] Yiwei Wang, Wei Wang, Yujun Ca, Bryan Hooi, Beng Chin Ooi. Detecting implementation bugs in graph convolutional network based node classifiers, *ISSRE 2020*.
- [108] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, Juncheng Liu, Bryan Hooi. NodeAug: Semi-supervised node classification with data augmentation, *KDD 2020*. (CCF-A)
- [109] Licheng Zhao, Yiwei Wang, Sandeep Kumar, Daniel P Palomar. Optimization algorithms for graph Laplacian estimation via ADMM and MM, *TSP 2019*. (JCR Q1)
- [110] Yiwei Wang, Mark James Carman, Yuan-Fang Li. Using knowledge graphs to explain entity co-occurrence in Twitter, *CIKM 2017*.

学术会议（高级）领域主席 ((Senior) Area Chair)

- Association for the Advancement of Artificial Intelligence (AAAI) (CCF-A)
- Conference on Neural Information Processing Systems (NeurIPS) (CCF-A)
- International Conference on Learning Representations (ICLR) (CCF-A)
- International Conference on Machine Learning (ICML) (CCF-A)
- Conference on Empirical Methods in Natural Language Processing (EMNLP)
- Annual Meeting of the Association for Computational Linguistics (ACL) (CCF-A)
- North American Chapter of the Association for Computational Linguistics (NAACL)
- Association for Computational Linguistics Rolling Review (ACL ARR)
- European Conference on Computer Vision (ECCV) (CCF-A)
- International Conference on Computer Vision (ICCV) (CCF-A)
- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (CCF-A)
- Conference on Language Modeling (COLM)
- International Joint Conference on Artificial Intelligence (IJCAI) (CCF-A)
- IEEE International Conference on Multimedia and Expo (ICME)
- International Joint Conference on Neural Networks (IJCNN)

学术期刊副编辑 (Journal Associate Editor)

- Transactions on Machine Learning Research (TMLR)
- IEEE Transactions on Industrial Informatics (IEEE TII) (JCR Q1)
- Journal of Visual Communication and Image Representation
- Pattern Recognition (JCR Q1)
- Neurocomputing (JCR Q1)

会议与期刊审稿服务

- 国际会议审稿人
EMNLP; ACL; NAACL; KDD; AISTATS; LoG; WWW; ICLR; NeurIPS; AAAI; CVPR; ICCV; ECCV; IJCAI; ICME; ECML-PKDD; IJCNN
- 国际期刊审稿人
IEEE TPAMI; IEEE TKDE; IET Image Processing; IEEE Transactions on Big Data; Pattern Recognition;

学术会议中主办 Tutorial

- | | |
|------------------------|-----------|
| • 从感知到模拟：多模态推理中世界模型的兴起 | ECCV 2026 |
| • 从感知到模拟：多模态推理中世界模型的兴起 | CVPR 2026 |
| • 面向视觉语言模型的综合推理 | ICCV 2025 |

主持经费

- | | |
|--|------|
| • U.S. Bank 学术基金 | 2026 |
| • 美国国家科学基金会 NSF CISE Research Initiation Initiative (NSF CRII) | 2025 |
| • NVIDIA 学术基金 | 2025 |
| • 加州大学默塞德分校 Faculty Research Award | 2025 |

奖项荣誉

- | | |
|-------------------------------------|------|
| • ICML AI for Science Workshop 最佳论文 | 2026 |
| • EMNLP 最佳论文 Finallist | 2024 |
| • 新加坡国立大学研究生院杰出研究奖（每年 3-4 名） | 2022 |
| • 新加坡国立大学计算学院研究成就奖 | 2020 |
| • 新加坡数据科学学会杰出博士奖学金（每年 10 名） | 2021 |
| • 中华人民共和国教育部国家奖学金（前 3%） | 2015 |

教学经历

- | | |
|-----------------------------------|-------------|
| • 加州大学默塞德分校, CSE 188 自然语言处理课程主讲教师 | 2026 |
| • 加州大学默塞德分校, EECS 224 大语言模型课程主讲教师 | 2025 |
| • 新加坡国立大学, 大数据系统课程助教 | 2021 |
| • 新加坡国立大学, 编程方法课程助教 | 2021 |
| • 新加坡国立大学, 知识发现与数据挖掘课程助教 | 2020、2021 年 |
| • 新加坡国立大学, 并行计算课程助教 | 2020 |
| • 香港科技大学, 信号处理与通信课程助教 | 2018 |