

Chaojie Wang

+65-83097904 | xd_silly@163.com | github.com/chaojiewang94 | linkedin.com/in/chaojie-wang-071256194

Personal Profile

I am currently working as a Research Scientist in Skywork AI, which is led by **Prof Shuicheng Yan** and focuses on accelerating the pace of realization of Artificial General Intelligence (AGI). Before that I obtained my Ph.D degree from Xidian University in 2021, supervised by **Prof Bo Chen** and **Prof Mingyuan Zhou**, and worked as Research Fellow in Nanyang Technological University until 2023, supervised by **Prof Bo An**.

My research interests mainly include generative models (Diffusion Models and Large Language Models), reinforcement learning from human feedback (RLHF) technologies and their applications on AIGC related tasks, such as retrieval-augmented generation (RAG), alignment in modality generation and transfer. Besides, I am also familiar with other Machine Learning related field, like Bayesian Models, Graph Neural Networks, Self-supervised Learning and Out-of-distribution Detection etc.

Education

Xidian University

Ph.D in Signal and Information Processing

Xian, Shannxi, China

Sept 2016 - June 2021

Xidian University

B.E in Electronic Engineering and Information Science

Xian, Shannxi, China

Sept 2012 - June 2016

Work Experience

Nanyang Technological University

Postdoc in School of Computer Science and Engineering

Singapore, Singapore

Nov 2021 - Feb 2024

Skywork AI

Research Scientist in Alignment and Agents

Singapore, Singapore

Feb 2024 - Present

Professional Skills

Programming Python (pytorch, numpy, scikit-learn, etc.), MATLAB, CUDA, C/C++

Miscellaneous Linux, Shell (Bash), $\text{\LaTeX}$ (Overleaf/Markdown), Microsoft Office (word, excel, powerpoint, visio), Git.

Languages Mandarin Chinese (native proficiency), English (CET-6).

Awards

2021–2022 L2RPN-ICAPS Learning to Run a Power Network Challenge, **Champion**, 1st Place

Nanjing, China

2015–2016 ACM-ICPC China Shaanxi Provincial Programming, **Bronze Medal**, 3rd Place

Xian, China

Recent Projects

Q*: Improving Multi-step Reasoning for LLMs with Deliberative Planning

Singapore, Singapore

Skywork AI

Feb 2024 - Present

- Formalize the multi-step reasoning of LLMs as a Markov Decision Process (MDP) and learn a plug-and-play Q-value model as heuristic function for estimating expected future rewards.
- Built upon A* search, we use Q-value model to guide LLMs to select the most promising next reasoning step without fine-tuning.

Removing Length Bias in RLHF Is Not Enough

Singapore, Singapore

Skywork AI

Nov 2023 - June 2024

- Reveal that prompt-template bias learned by RMs can also cause reward hacking when dealing with marginal samples, resulting in LLMs preferring to generate responses in a specific format after RLHF fine-tuning.
- Propose a low-cost but effective method, namely Prompt Bias Calibration (PBC), to estimate the prompt-template bias term during reward modeling, which can be utilized to calibrate reward scores in the following RL fine-tuning process.

Large Langue Models Retrieved on Knowledge Graph

Singapore, Singapore

Nanyang Technological University

Dec 2022 - June 2023

- Develop a novel framework to accomplish KBQA tasks with LLMs, whose workflow mainly consists of four stages, specifically *Question Decomposition*, *Knowledge Retrieval*, *Candidate Reasoning*, and *Response Generation*, greatly improving the reliability of existing LLM's response.
- Moving beyond text-to-SQL generation, we introduce question templates as an intermediary to make it easy for LLMs to capture the logic of question decomposition, where each question template can be solved with several pre-collected logical chains.

Diffusion Models Augmented with Episodic-Semantic Memory (NeurIPS 2023)

Singapore, Singapore

Nanyang Technological University

Dec 2021 - Feb 2023

- Design a variational structured memory module (VSM), which can simultaneously store both episodic and semantic memories to assist existing generative models efficiently recall these memories during sample generation.
- Introduce a bionic memory updating strategy for the conversion between episodic and semantic memories, which can also model the uncertainty during conversion.

- Construct a variational graph autoencoder to infer the probability that each node belongs to various communities.
- Unify the generation of node edges and features under the same generative framework and interpret the underlying mechanism of community formation.
- Combine the procedure of community detection and the training of downstream tasks, which can provide more expressive node representations for graph analytic task, like node/graph classification.

Academic Services

- I have consistently served as a dedicated reviewer for AAAI 2021, ICLR 2021-2025, ICML 2021-2025, NeurIPS 2022-2025, IEEE TPAMI, IEEE TSP, IEEE TNNLS.

Selected Publications

First Author, Equal Contribution* and Corresponding Author#

Alignment and Agents Related

- Jiakai Liu, **Chaojie Wang#**, and Yang Liu. Improving multi-step reasoning abilities of large language models with direct advantage policy optimization [C]. **NeurIPS Spotlight**, 2025.
- **Chaojie Wang***, Yancheng Deng*#, Zhiyi Lv, Shuicheng Yan, and Bo An. Q*: Improving Multi-step Reasoning for LLMs with Deliberative Planning [C]. TPAMI, 2025 (Under Review).
- **Chaojie Wang**, Haonan Shi, Long Tian#, Bo An, and Shuicheng Yan. Removing Length Bias in RLHF Is Not Enough [C], ACL, 2025.
- Chuqiao Zong, **Chaojie Wang#**, Molei Qin, Lei Feng, Xinrun Wang, and Bo An. MacroHFT: Memory Augmented Context-aware Reinforcement Learning On High Frequency Trading [C]. KDD, 2024.

Natural Language Processing Related

- **Chaojie Wang**, Bo Chen, Zhibin Duan, Wenchao Chen, Hao Zhang, and Mingyuan Zhou. Generative Text Convolutional Neural Network for Hierarchical Document Representation Learning [J]. TPAMI, 2022.
- Hao Zhang*, **Chaojie Wang***, Zhengjue Wang, Zhibin Duan, Bo Chen, Mingyuan Zhou, Ricardo Henao, and Lawrence Carin. Learning Hierarchical Documents Graphs From Multilevel Sentence Relations [J]. TNNLS, 2022.
- Wenchao Chen*, **Chaojie Wang***, Bo Chen, Yicheng Liu, Hao Zhang, and Mingyuan Zhou. Bidirectional Convolutional Poisson Gamma Dynamical Systems [C]. NeurIPS, 2020.
- **Chaojie Wang**, Bo Chen, Sucheng Xiao, and Mingyuan Zhou. Convolutional Poisson Gamma Belief Network [C]. ICML, 2019.

Generative Model Theory Related

- Shuo He, **Chaojie Wang**, Guowu Yang, and Feng Lei. Candidate Label Set Pruning: A Data-centric Perspective for Deep Partial-label Learning [C], **ICLR Oral**, 2024.
- Zhibin Duan*, Zhiyi Lv*, **Chaojie Wang#**, Bo Chen, Bo An, and Mingyuan Zhou. Few-shot Generation via Recalling the Episodic-Semantic Memory like Human Being [C], NeurIPS, 2023.
- Yewen Li*, **Chaojie Wang*#**, Xiaobo Xia, Tongliang Liu, Xin Miao, and Bo An. Out-of-Distribution Detection with An Adaptive Likelihood Ratio on Informative Hierarchical VAE [C]. NeurIPS, 2022.
- Yewen Li*, **Chaojie Wang*#**, Zhibin Duan, Dongsheng Wang, Bo Chen, Bo An, and Mingyuan Zhou. Alleviating “Posterior Collapse” in Deep Topic Models via Policy Gradient [C]. NeurIPS, 2022.

Graph Neural Networks Related

- Yilin He*, **Chaojie Wang*#**, Hao Zhang, Bo Chen, and Mingyuan Zhou#. A Variational Edge Partition Model for Supervised Graph Representation Learning [C]. NeurIPS, 2022.
- **Chaojie Wang***, Hao Zhang*, Bo Chen, Dongsheng Wang, Zhengjue Wang, and Mingyuan Zhou. Deep Relational Topic Modeling via Graph Poisson Gamma Belief Network [C]. NeurIPS, 2020.
- Zhengjue Wang*, **Chaojie Wang***, Hao Zhang, Zhibin Duan, Bo Chen, and Mingyuan Zhou. Learning Dynamic Hierarchical Topic Graph with Graph Convolutional Network for Document Classification [C]. AISTATS, 2020.

Multi-Modality Related

- **Chaojie Wang**, Bo Chen, Sucheng Xiao, Zhengjue Wang, Hao Zhang, Penghui Wang, Han Ning, and Mingyuan Zhou. Multimodal Weibull Variational Autoencoder for Jointly Modeling Image-Text Data [J]. TCYB 2021.
- **Chaojie Wang**, Bo Chen, and Mingyuan Zhou. Multimodal Poisson Gamma Belief Network [C]. AAAI, 2018.