

Peifeng Wang

AI Research Scientist, Meta, US

📞 +1 2132040743 • ✉ peifengw@usc.edu • 🌐 wangpf3.github.io

Research Interests

- Reinforcement Learning for Reasoning, Large Language Models

Employment History

Meta FAIR

Research Scientist

Menlo Park, USA

Apr 2025 - now

- Focus: LLM Reasoning
- **RL in Unverifiable Domains**: RL with rubric-based rewards in practical domains (e.g., finance, health)
- **Mid-Training for RL**: Integrate thinking traces in mid-training to build a strong base model for reasoning
- **Scaling RL**: Build efficient and scalable RL recipe for long-context reasoning

Salesforce AI

Research Scientist

Palo Alto, USA

Jan 2024 - Apr 2025

- Focus: LLM Reasoning, LLM Agent, AI feedback for model development
- LLM-as-Judge: Build a generative judge model that generalizes across various evaluation tasks [Blog][Paper]
- CRM Benchmarking: Develop the world's first LLM benchmark for CRM to assess the efficacy of generative AI models for business applications [Blog][Leaderboard]

Meta FAIR

Research Intern

New York, USA

May 2023 - Aug. 2023

- Topic: Multimodal Chain-of-Thought Reasoning
- Managers: Dr. Maryam Fazel-Zarandi, Dr. Asli Celikyilmaz

Amazon

Applied Scientist Intern

Palo Alto, USA

May 2022 - Nov. 2022

- Topic: Faithful Chain-of-Thought Reasoning
- Managers: Dr. Zhengyang Wang, Dr. Zheng Li

Microsoft

Research Intern

Remote, USA

May 2021 - Aug. 2021

- Topic: Knowledge-Enhanced Large Scale Language Model Pre-training
- Manager: Dr. Dinghan Shen

Tencent AI Lab

Research Intern, NLP Group

Shenzhen, China

Jun. 2018 - Sep. 2018

- Topic: Inductive Knowledge Representation Learning
- Manager: Dr. Jialong Han

Academic History

University of Southern California, Los Angeles, US

- Advisor: Prof. Xiang Ren, Prof. Muhao Chen,

Ph.D. in Computer Science

Aug. 2019 - Dec. 2023

Sun Yat-sen University, Guangzhou, China

- Advisor: Prof. Rong Pan

M.S. in Computer Science

Sep. 2016 - Jun. 2019

Sun Yat-sen University, Guangzhou, China

- GPA: 3.8/4.0

B.S. in Computer Science

Sep. 2012 - Jun. 2016

Selected Publications (Google Scholar Profile)

Evaluating judges as evaluators: The benchmark of llm-judges as test-time scaling evaluators

- Yilun Zhou, Austin Xu, **Peifeng Wang**, Caiming Xiong, Shafiq Joty
- *ICML 2025*

A survey of frontiers in llm reasoning: Inference scaling, learning to reason, and agentic systems

- Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, **Peifeng Wang**, Silvio Savarese, Caiming Xiong, Shafiq Joty
- *TMLR 2025*

Direct Judgement Preference Optimization

- **Peifeng Wang**, Austin Xu, Yilun Zhou, Caiming Xiong, Shafiq Joty
- *EMNLP 2025*

DOMINO: A Dual-System for Multi-step Visual Language Reasoning

- **Peifeng Wang**, Olga Golovneva, Armen Aghajanyan, Xiang Ren, Muhao Chen, Asli Celikyilmaz, Maryam Fazel-Zarandi
- *arXiv*

SCOTT: Self-Consistent Chain-of-Thought Distillation

- **Peifeng Wang**, Zhengyang Wang, Zheng Li, Yifan Gao, Bing Yin, Xiang Ren
- *ACL 2023 Outstanding Paper Award*

PINTO: Faithful Language Reasoning Using Prompt-Generated Rationales

- **Peifeng Wang**, Aaron Chan, Filip Ilievski, Muhao Chen, Xiang Ren
- *ICLR 2023. A short version is accepted by TL4NLP and TSRML Workshop at NeurIPS, 2022*

Contextualized Scene Imagination for Generative Commonsense Reasoning

- **Peifeng Wang**, Jonathan Zamora, Junfeng Liu, Filip Ilievski, Muhao Chen, Xiang Ren
- *ICLR 2022*

Learning to Deceive Knowledge Graph Augmented Models via Targeted Perturbation

- Mrigank Raman, Siddhant Agarwal, **Peifeng Wang**, Aaron Chan, Hansen Wang, Sungchul Kim, Ryan Rossi, Handong Zhao, Nedim Lipka, Xiang Ren
- *ICLR 2021*

Connecting the Dots: A Knowledgeable Path Generator for Commonsense Question Answering

- **Peifeng Wang**, Nanyun Peng, Filip Ilievski, Pedro Szekely, Xiang Ren
- *EMNLP 2020 (Findings)*

Scalable Multi-Hop Relational Reasoning for Knowledge-Aware Question Answering

- Yanlin Feng, Xinyue Chen, Bill Yuchen Lin, **Peifeng Wang**, Jun Yan, Xiang Ren
- *EMNLP 2020*

Mentoring

- Jonathan Zamora (2022), Undergraduate Student at UCSD
- Junfeng Liu (2022), Master's Student at USC
- Mrigank Raman (2021), Undergraduate Student at IIT Delhi

Awards

- ACL Outstanding Paper Award, 2023
- Annenberg Fellowship Award, University of Southern California, 2019.