ZHE WANG

zw6sg@virginia.edu, Arlington, Virginia, United States, LinkedIn

PROFESSIONAL EXPERIENCE

Amazon, United States *Jun. 2024 – Present*

Applied Scientist II

LLM Post-Training and Model Customization

- Core developer of Amazon Bedrock and SageMaker LLM post-training services, contributing to supervised fine-tuning (SFT), preference fine-tuning (PFT), reinforcement fine-tuning (RFT), and multi-turn training workflows.
- Key technical contributor to the end-to-end launch of RLVR (Reinforcement Learning with Verifiable Rewards) in Amazon SageMaker, enabling reinforcement learning workflows for foundation models.
- Led the implementation and productionization of fine-tuning support for the NVIDIA Nemotron family across serverful and serverless environments, partnering with SageMaker engineering teams for end-to-end integration.
- Partnered with Salesforce to develop AgentCore fine-tuning workflows, highlighted during AWS re:Invent 2026.
- Built the official LLM fine-tuning cookbook covering SFT, RL, tool use, function calling, and conversational agent applications to accelerate customer adoption.
- Co-organized the KDD 2026 hands-on tutorial: *Multi-Turn Reinforcement Learning for Large Language Models: From Theory to Practice with Amazon SageMaker AI*.

LLM Safety and Automated Red Teaming

- Developed automated red-teaming approaches for large language models and multimodal language models; integrated safety testing workflows into Amazon Bedrock Guardrails for synthetic safety data generation.
- Designed scalable evaluation pipelines to benchmark safety behaviors of 20+ commercial and open-source LLMs and multimodal LLMs.
- Built an extensible stress-testing framework supporting state-of-the-art jailbreak attacks for continuous model evaluation.

Multimodal Model Evaluation

- Developed datasets, evaluation dimensions, metrics, and automated evaluators for benchmarking text, image, and video generation models, including tuning multimodal LLMs as automated judges.

The University of Virginia, United States *Aug. 2019 – May. 2024*

Research Assistant

Generalization of the deep learning models

- Developed robust and reliable deep neural networks to generalize to unseen data, tasks, distributions, and domains. This contributes to robust AI services for time-series forecasting, classification, regression, and generation.

INTERN EXPERIENCE

Applied Scientist Intern, AWS Generative AI Innovation Center, Amazon *May. 2023 – Aug. 2023*

- Developed adversarial suffix-based jailbreak methods to evaluate and improve LLM safety. Published resulting work at NAACL 2025.

Applied Scientist Intern, Machine Learning Solution Lab, Amazon *May. 2022 – Aug. 2022*

- Developed retrieval-augmented transformer models for few-shot offline reinforcement learning.
- Achieved state-of-the-art performance across seven robotic manipulation benchmarks. Published resulting work at WWW 2025.

EDUCATION

The University of Virginia, United States *Aug. 2019 – May 2024*

Ph.D. in Computer Science
Advisor: Dr. Yanjun Qi

Research Areas: Large Language Model Safety, Adversarial Machine Learning, Reinforcement Learning, Meta-Learning, Domain Generalization, Multimodal AI.

Xi'an Jiaotong University, China *Aug. 2016 – Jun. 2019*

Master of Applied Mathematics
Advisor: Dr. Jian Sun

Research Areas: Optimization acceleration and deep learning architectures for low-rank constrained optimization.

Xi'an Jiaotong University, China *Aug. 2012 – Jun. 2016*

Bachelor of Mathematics (Honors Program)

SELECTED PUBLICATIONS

1. **Zhe Wang**, Sapana Chaudhary, Jiayu Li, Yuyang Wang, Xuan Zhu. *Multi-Turn Reinforcement Learning for Large Language Models: From Theory to Practice with Amazon SageMaker AI*, KDD 2026 Hands-on Tutorials.
2. Shiyu Yuan, Sourav Sanjukta Bhabesh, **Zhe Wang**, Dmitriy Bespalov, Wesley Rose, Huzefa Rangwala. *SD-MAR: Multi-image Analytical Reasoning via Synthetic Data and Reinforcement Learning* , Arxiv 2026.
3. Yidong Ouyang, **Zhe Wang**, Sourav Bhabesh, Dmitriy Bespalov. *Reinforcing the Generation Order of Multimodal Masked Diffusion Models*, Arxiv 2026.
4. Yidong Ouyang, Panwen Hu, Zhengyan Wan, **Zhe Wang**, Liyan Xie, Dmitriy Bespalov, Ying Nian Wu, Guang Cheng, Hongyuan Zha, Qiang Sun. *Training-Free Self-Correction for Multimodal Masked Diffusion Models*, Under review, 2026.
5. **Zhe Wang**, Yanjun Qi. *Augmented Adversarial Trigger Learning*, NAACL 2025.
6. Aman Goel, Xian Wu, **Zhe Wang**, Dmitriy Bespalov, Yanjun Qi. *TurboFuzzLLM: Turbocharging Mutation-based Fuzzing for Effectively Jailbreaking Large Language Models in Practice*, NAACL 2025, **Oral Acceptance**.
7. **Zhe Wang**, Haozhu Wang, Dmitriy Bespalov, and Yanjun Qi. *T3GDT: Three Tier Tokens Guiding Decision Transformer for Offline Reinforcement Learning*, WWW, 2025.
8. Daniel Schwartz, Dmitriy Bespalov, **Zhe Wang**, Ninad Kulkarni, Yanjun Qi. *Automated Generation of Stealthy Jailbreak Prompts and its Application to Improving Content Moderation*, EMNLP, 2025.
9. **Zhe Wang**, Jake Grigsby, and Yanjun Qi. *PGrad: Learning robust gradients for domain generalization.* ICLR, 2023.
10. **Zhe Wang**, Jake Grigsby, Arshdeep Sekhon, and Yanjun Qi. *ST-MAML: A Stochastic-Task based Method for Task-Heterogeneous Meta-Learning.* UAI, 2022.
11. Jake Grigsby, **Zhe Wang**, and Yanjun Qi. *Long-Range Transformers for Dynamic Spatiotemporal Forecasting.* arXiv preprint arXiv: 2109.12218, 2021. ($>$ 800 **stars in GitHub**)
12. Arshdeep Sekhon, Hanjie Chen, **Zhe Wang**, Yangfeng Ji, and Yanjun Qi. *Learning Variational Word Interaction Graph for Interpretable Neural Text Classifier.* AAAI, 2023, **Oral acceptance**.

SERVICE

Reviewer: ICLR, ICML, NeurIPS, CVPR, AAAI, IJCAI, AISTATS, UAI (2024–2026)

Teaching Assistant: Machine Learning, Statistical Learning, University of Virginia

TECHNICAL SKILLS

Generative AI: LLMs, Multimodal LLMs, SFT, RLHF, RLVR, DPO, PEFT, LoRA, Model Evaluation

Machine Learning: Deep Learning, Adversarial ML, Meta-Learning, Domain Generalization

Frameworks and Systems: PyTorch, Transformers, DeepSpeed, Megatron-LM, SageMaker, Amazon Bedrock