

KUN ZHAO

Ph.D. Candidate, Electrical & Computer Engineering, University of Pittsburgh

kun.zhao@pitt.edu • (+1) 412-519-4341 • Pittsburgh, PA, USA

Homepage • Google Scholar • GitHub

Research Interests

LLM Post-Training & Reinforcement Learning • Multimodal Reasoning • Reliable Evaluation & Alignment • Efficient LLM Systems

Ph.D. candidate at the University of Pittsburgh developing methods for LLM post-training, multimodal reasoning, and reliable evaluation. My research focuses on reinforcement-learning-based alignment, fine-grained reward design, and robust evaluation of open-ended generation. First/co-first author of ACL and EMNLP papers, including an ACL main-conference oral, with hands-on experience building distributed LLM training and online rollout systems.

Selected Publications

- **HERO: Hierarchical Evidential Reasoning Optimization for Radiology Report Generation via Reason-then-Summarize**

Kun Zhao, Guodong Liu, Hui Ji, Siyuan Dai, Pan Wang, Jifeng Song, Chenghua Lin, Liang Zhan, Haoteng Tang

arXiv:2601.03321 • Under review, AAAI 2027 • First author

Contribution: Developed factorized GRPO credit assignment with segment-specific advantages, pathology-aware token weighting, and an evidence-grounded completion gate for radiology report generation.

- **Evaluating Open-Domain Dialogues in Latent Space with Next Sentence Prediction and Mutual Information**

Kun Zhao, Bohao Yang, Chenghua Lin, Wenge Rong, Aline Villavicencio, Xiaohui Cui

ACL 2023 Main Conference • Oral presentation • First author

Contribution: Augmented a conditional VAE with next-sentence prediction and mutual information to evaluate one-to-many dialogue responses in latent space, including valid responses semantically distant from a single reference.

- **X-ray Made Simple: Lay Radiology Report Generation and Robust Evaluation**

Kun Zhao, Chenghao Xiao, Sixing Yan, Haoteng Tang, William K. Cheung, Noura Al Moubayed, Liang Zhan, Chenghua Lin

ACL 2026 Findings • First author

Contribution: Created sentence- and report-level lay radiology datasets, a semantics-based metric that exposes template-driven BLEU inflation, and a lay-guided training strategy.

Education

University of Pittsburgh

Ph.D. in ECE • GPA 3.925/4.00 • Advisor: Prof. Liang Zhan

Pittsburgh, PA, USA

Sep. 2023 – Expected 2027

Wuhan University

M.Eng. in Cyberspace Security, School of Cyber Science and Engineering

Wuhan, China

Sep. 2019 – Jun. 2022

Zhengzhou University of Light Industry

B.Eng. in Internet of Things Engineering

Zhengzhou, China

Sep. 2014 – Jun. 2018

Research Experience

University of Pittsburgh — Dept. of Electrical & Computer Engineering

Pittsburgh, PA, USA

Graduate Research Assistant · Advisor: Prof. Liang Zhan

Sep. 2023 – present

- Lead author and project lead on LLM post-training, reliable evaluation of open-ended generation, and multimodal learning, with medical AI as an application domain.

Reinforcement Learning for Faithful Multimodal Generation

2024 – present

Project lead · with Prof. Liang Zhan, Prof. Chenghua Lin, Prof. Haoteng Tang

- Developed **HERO** on Qwen-2.5-VL-7B with a two-part output: free-text radiological findings followed by a 14-label CheXpert diagnostic JSON.
- Replaced vanilla GRPO's broadcast advantage with **segment-specific advantages**, TF-IDF and pathology-aware token weighting, and an **evidence-grounded completion gate**; designed a nine-component reward bank spanning reasoning, consistency, grounding, validity, and diagnostic correctness.
- On MIMIC-CXR, achieved RadGraph 0.354 and report-level Micro-F1(14) 0.574; reached think-answer SCS 0.980 and evidence grounding 0.773, with transfer evaluation on IU-Xray.
- Implemented two-stage SFT followed by GRPO with vLLM rollouts, per-token importance-sampling correction, KL regularization, and conditional ground-truth injection on NVIDIA A100 GPUs.

Reliable Evaluation of Open-Ended Generation

2021 – present

Project lead · with Prof. Chenghua Lin, Prof. Liang Zhan

- Developed a line of **reference-free evaluation methods**, progressing from latent-space response modeling (**CMN**) to SLM-LLM collaboration (**SLIDE/DRE**) and structured semantic augmentation with **AMR**.
- Introduced CMN by augmenting a conditional VAE with next-sentence prediction and mutual information, improving evaluation of valid dialogue responses that are semantically distant from a single reference (ACL 2023 Main, **Oral; first author**).
- Developed SLIDE: a contrastively trained SLM separates robust response features and supplies a semantic-sensitivity score that is adaptively combined with an LLM evaluator (ACL 2024 Findings). DRE adds SLM-guided interior prompt refinement and exterior score refinement, improving alignment with human judgments across dialogue benchmarks (*Neural Networks*, 2026, **in press**).
- Co-developed an AMR-enhanced evaluator that gates AMR graph information into an SLM and injects both AMR knowledge and SLM predictions into LLM prompts, improving discrimination of lexically similar adversarial responses (EMNLP 2025 Findings, **co-first author**).

Multimodal Foundation Models and Representation Learning

2024 – present

Project lead · with Prof. Liang Zhan, Prof. Haoteng Tang

- Developed **R-GenIMA**, which represents parcellated 3D MRI regions as visual tokens and SNP profiles as prompted genetic tokens for four-stage Alzheimer's disease classification; ROI and gene attribution analyses recover stable, biologically plausible patterns (*Frontiers in Radiology*, 2026, first author).
- Created refined sentence- and report-level lay radiology datasets and a semantics-based metric that exposes template-driven BLEU inflation; showed that lay-guided training improves semantic learning as data scale increases (ACL 2026 Findings, first author). The data were used in the BioLay-Summ 2025 radiology-report track.
- Systematically benchmarked vision-encoder and language-backbone combinations for radiology report generation through controlled component swaps (ICDM Workshops 2025, first author).
- Contributed to STE-ODE for structural-effective brain-network dynamics (MICCAI 2024) and to FigEx2, which localizes panels and generates panel-level captions directly from compound scientific figures (EMNLP 2026 Main Conference).

Publications

Grouped by publication type; dated papers in reverse chronological order. **Bold** indicates my name; * indicates equal contribution.

Peer-Reviewed Conference Papers

- [1] FigEx2: Visual-Conditioned Panel Detection and Captioning for Scientific Compound Figures
Jifeng Song, Arun Das, Pan Wang, Hui Ji, **Kun Zhao**, Yufei Huang
EMNLP 2026 Main Conference · Accepted
- [2] X-ray Made Simple: Lay Radiology Report Generation and Robust Evaluation
Kun Zhao, Chenghao Xiao, Sixing Yan, Haoteng Tang, William K. Cheung, Noura Al Moubayed, Liang Zhan, Chenghua Lin
ACL 2026 Findings | **First author**
- [3] Emphasising Structured Information: Integrating Abstract Meaning Representation into LLMs for Enhanced Open-Domain Dialogue Evaluation
Bohao Yang*, **Kun Zhao***, et al.
EMNLP 2025 Findings | **Co-first author**
- [4] Crafting Customisable Characters with LLMs: A Persona-Driven Role-Playing Agent Framework
Bohao Yang, Dong Liu, Chenghao Xiao, **Kun Zhao**, Chen Tang, Chao Li, Lin Yuan, Yang Guang, Chenghua Lin
EMNLP 2025 Findings
- [5] Who Matters More in Radiology Report Generation: Vision Encoders or Language Models?
Kun Zhao, Yang Du, Liang Zhan, Haoteng Tang, et al.
SAIMBio Workshop @ ICDM 2025 | **First author**
- [6] Interpretable Spatio-Temporal Embedding for Brain Structural-Effective Network with Ordinary Differential Equation
Haoteng Tang, Guodong Liu, Siyuan Dai, Kai Ye, **Kun Zhao**, Wenlu Wang, Carl Yang, Lifang He, Alex Leow, Paul Thompson
MICCAI 2024
- [7] SLIDE: A Framework Integrating Small and Large Language Models for Open-Domain Dialogues Evaluation
Kun Zhao, Bohao Yang, Chen Tang, Chenghua Lin, Liang Zhan
ACL 2024 Findings | **First author**
- [8] Effective Distillation of Table-Based Reasoning Ability from LLMs
Bohao Yang, Chen Tang, **Kun Zhao**, Chenghao Xiao, Chenghua Lin
LREC-COLING 2024
- [9] Evaluating Open-Domain Dialogues in Latent Space with Next Sentence Prediction and Mutual Information
Kun Zhao, Bohao Yang, Chenghua Lin, Wenge Rong, Aline Villavicencio, Xiaohui Cui
ACL 2023 Main Conference | **First author · Oral presentation**
- [10] A Latent Variable Model with Hierarchical Structure and GPT-2 for Long Text Generation
Kun Zhao, Hongwei Ding, Kai Ye, Xiaohui Cui, Zhongwang Fu
International Conference on Artificial Neural Networks (ICANN) 2021 | **First author**

Journal Articles

- [11] DRE: An Effective Dual-Refined Method for Integrating Small and Large Language Models in Open-Domain Dialogue Evaluation
Kun Zhao, Bohao Yang, Chen Tang, et al.
Neural Networks, 2026, 109574 · *In press; available online* | **First author**
- [12] Interpretable Multimodal Learning for Integrating Neuroimaging and Genetic Data in Alzheimer’s Disease
Kun Zhao, Siyuan Dai, Yingying Zhang, Guodong Liu, Pengfei Gu, Chenghua Lin, Paul M. Thompson, Alex Leow, Heng Huang, Lifang He, Liang Zhan, Haoteng Tang
Frontiers in Radiology, 2026 | **First author**
- [13] Constrained Multiview Contrastive Learning for Jointly Supervised Representation Learning
Siyuan Dai, Kai Ye, **Kun Zhao**, Yang Du, Haoteng Tang, Liang Zhan
Meta-Radiology, 4(3):100216, 2026
- [14] A Transformer-Based Hierarchical Variational AutoEncoder Combined Hidden Markov Model for Long Text Generation
Kun Zhao, Hongwei Ding, Kai Ye, Xiaohui Cui
Entropy, 23(10):1277, 2021 | **First author**

Manuscripts Under Review

- [15] HERO: Hierarchical Evidential Reasoning Optimization for Radiology Report Generation via Reason-then-Summarize
Kun Zhao, Guodong Liu, Hui Ji, Siyuan Dai, Pan Wang, Jifeng Song, Chenghua Lin, Liang Zhan, Haoteng Tang
arXiv:2601.03321 · *Under review, AAAI 2027* | **First author**

Academic Service

Organizer, BioLaySumm Shared Task, BioNLP Workshop @ ACL 2025

Task design, benchmark construction, and evaluation protocol (biolaysumm.org).

Reviewer, ACL Rolling Review (ARR), NeurIPS, AAAI

Technical Skills

Languages	Python, C/C++, Bash, $\text{\LaTeX}$
LLM/ML	PyTorch, Hugging Face Transformers, TRL, DeepSpeed (ZeRO-3), vLLM, PEFT/LoRA, LLaVA, LLaMA-Factory
Methods	SFT, RLHF/GRPO, fine-grained reward design, multimodal alignment, LLM-as-a-judge and reference-free evaluation, distributed training and inference
Systems	Linux, Slurm, Docker, Git