

KUN ZHAO

Ph.D. Candidate, Electrical & Computer Engineering, University of Pittsburgh
Reliable Evaluation · LLM Post-Training · Multimodal Medical AI
kun.zhao@pitt.edu • (+1) 412-519-4341 • Pittsburgh, PA, USA
Homepage • Google Scholar • GitHub

Research Interests

Reliable Evaluation of Open-Ended Generation · LLM Post-Training & Reinforcement Learning · Multimodal Reasoning · Medical AI

My research asks how reliable evaluation can guide better learning in generative AI. I study text-only evaluation of open-domain dialogue, structured feedback for reinforcement-learning post-training, and evidence-grounded multimodal generation and interpretable medical AI. Across these settings, I design evaluators and learning signals that reflect the structure of generated outputs and the evidence available for each task.

Selected Research Contributions

Representative first- and co-first-author work. Complete record under *Publications*. * denotes equal contribution.

- **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: Introduced CMN, which augments a conditional VAE with next-sentence prediction and mutual information to evaluate one-to-many dialogue responses in latent space.

- **DRE: An Effective Dual-Refined Method for Integrating Small and Large Language Models in Open-Domain Dialogue Evaluation**

Kun Zhao, Bohao Yang, Chen Tang, Siyuan Dai, Haoteng Tang, Chenghua Lin, Liang Zhan

Neural Networks, 2026, 109574 · First author

Contribution: Combined SLM-guided interior prompt refinement with exterior score refinement to improve alignment with human judgments across open-domain dialogue benchmarks.

- **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 evaluation metric, and a lay-guided training strategy.

- **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.

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

Graduate Research Assistant · Advisor: Prof. Liang Zhan

Pittsburgh, PA, USA

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.

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, 109574; **first author**).
- 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**).

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.

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).

Wuhan University — Big Data and Artificial Intelligence Laboratory

Wuhan, China

Research Assistant · hierarchical latent-variable models for long-form text

Sep. 2022 – Apr. 2023

Publications

Grouped by publication type; first- and co-first-author work is listed before collaborative work. **Bold** indicates my name; * indicates equal contribution. Preprints superseded by a peer-reviewed version are consolidated under the final publication.

Peer-Reviewed Conference and Workshop Papers

- [1] 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**
- [2] Emphasising Structured Information: Integrating Abstract Meaning Representation into LLMs for Enhanced Open-Domain Dialogue Evaluation
Bohao Yang*, **Kun Zhao***, Dong Liu, Chen Tang, Liang Zhan, Chenghua Lin
EMNLP 2025 Findings | **Co-first author**
- [3] Who Matters More in Radiology Report Generation: Vision Encoders or Language Models?
Kun Zhao, Yang Du, Rhianna Zhang, Liang Zhan, Dongkuan Xu, Pengfei Gu, Haoteng Tang
SAIMBio Workshop @ IEEE ICDM Workshops 2025 | **First author**
- [4] 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**
- [5] 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**
- [6] 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**
- [7] 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

- [8] Overview of the BioLaySumm 2025 Shared Task on Lay Summarization of Biomedical Research Articles and Radiology Reports
Chenghao Xiao, **Kun Zhao**, Xiao Wang, Siwei Wu, Sixing Yan, Tomas Goldsack, Sophia Ananiadou, Noura Al Moubayed, Liang Zhan, William K. Cheung, Chenghua Lin
BioNLP Workshop @ ACL 2025, pp. 365–377
- [9] A General Paradigm for Fine-Tuning Large Language Models in Alzheimer’s Disease Diagnosis
Marcus Zhan, **Kun Zhao**, Guodong Liu, Haoteng Tang
AAAI Symposium Series, 5(1):37–42, 2025
- [10] Pretraining a Generative AI Model on CBT Sessions to Emulate Human Therapist Strategies: A Study Using the APA PsycTherapy Database
Fahad Islam, **Kun Zhao**, Bart P. Knijnenburg, Ming-Hui Chen, Alex Leow, Prasanna Muthukanagaraj, Liang Zhan, Chao Shi
SAIMBio Workshop @ IEEE ICDM Workshops 2025, pp. 2011–2014
- [11] Why Text Prevails: Vision May Undermine Multimodal Medical Decision Making
Siyuan Dai, Lunxiao Li, **Kun Zhao**, Eardi Lila, Paul K. Crane, Heng Huang, Dongkuan Xu, Haoteng Tang, Liang Zhan
IEEE ICDM Workshops 2025, pp. 2020–2024
- [12] 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
- [13] 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
- [14] Effective Distillation of Table-Based Reasoning Ability from LLMs
Bohao Yang, Chen Tang, **Kun Zhao**, Chenghao Xiao, Chenghua Lin
LREC-COLING 2024
- [15] Graph Enhanced BERT for Stance-Aware Rumor Verification on Social Media
Kai Ye, Yangheran Piao, **Kun Zhao**, Xiaohui Cui
International Conference on Artificial Neural Networks (ICANN) 2021, pp. 422–435

Journal Articles

- [16] DRE: An Effective Dual-Refined Method for Integrating Small and Large Language Models in Open-Domain Dialogue Evaluation
Kun Zhao, Bohao Yang, Chen Tang, Siyuan Dai, Haoteng Tang, Chenghua Lin, Liang Zhan
Neural Networks, 2026, 109574 | **First author**
- [17] 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**
- [18] 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**
- [19] 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

- [20] Brain Structural Network Alterations in Alzheimer’s Disease: Perspectives from Ethnicity-Specific Factors
Haoteng Tang, **Kun Zhao**, Guodong Liu, Paul M. Thompson, Heng Huang, Alex Leow, Liang Zhan
Alzheimer’s & Dementia, 21:e101821, 2025
- [21] A Heterogeneous Graph Enhanced LSTM Network for Hog Price Prediction Using Online Discussion
Kai Ye, Yangheran Piao, **Kun Zhao**, Xiaohui Cui
Agriculture, 11(4):359, 2021
- [22] MRU-Net: A U-Shaped Network for Retinal Vessel Segmentation
Hongwei Ding, Xiaohui Cui, Leiyang Chen, **Kun Zhao**
Applied Sciences, 10(19):6823, 2020

Manuscripts Under Review

- [23] 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**

Preprints

- [24] Longitudinal Bayesian Learning of Continuous Disease Position across the Alzheimer’s Disease Continuum
Yingying Zhang, **Kun Zhao**, Guodong Liu, Qi Huang, Pengfei Gu, Dongchul Kim, Erik Enriquez, Alex D. Leow, Paul M. Thompson, Heng Huang, Hongchang Gao, Liang Zhan, Haoteng Tang
arXiv:2608.19436, 2026
- [25] Artemis: Anatomy-Resolved inTervention for Eliminating Multimodal NeuroImage confounderS
Siyuan Dai, Yang Du, **Kun Zhao**, Zhusuyi Chen, Heng Huang, Paul Thompson, Chao Shi, Haoteng Tang, Liang Zhan
arXiv:2606.18287, 2026
- [26] GRAPHMOE: Amplifying Cognitive Depth of Mixture-of-Experts Network via Introducing Self-Rethinking Mechanism
Chen Tang, Bo Lv, Zifan Zheng, Bohao Yang, **Kun Zhao**, Ning Liao, Xiaoxing Wang, Feiyu Xiong, Zhiyu Li, Nanyu Liu, Jingchi Jiang
arXiv:2501.07890, 2025
- [27] BioMNER: A Dataset for Biomedical Method Entity Recognition
Chen Tang, Bohao Yang, **Kun Zhao**, Bo Lv, Chenghao Xiao, Frank Guerin, Chenghua Lin
arXiv:2406.20038, 2024
- [28] Constrained Multiview Representation for Self-Supervised Contrastive Learning
Siyuan Dai, Kai Ye, **Kun Zhao**, Ge Cui, Haoteng Tang, Liang Zhan
arXiv:2402.03456, 2024

Presentations

- **Poster presentation**, ACL 2026 — *X-ray Made Simple: Lay Radiology Report Generation and Robust Evaluation*.
- **Oral presentation**, ACL 2023 Main Conference, Toronto, Canada — *Evaluating Open-Domain Dialogues in Latent Space with Next Sentence Prediction and Mutual Information*.

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

Mentoring

- Mentored two cohorts of high school students in CoSBBI 2025 and CoSBBI 2026 through counseling-dialogue generation projects.

Honors and Awards

- Third-Class Academic Scholarship, Wuhan University 2020

Technical Skills

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

References

Prof. Liang Zhan	Doctoral advisor, University of Pittsburgh · liang.zhan@pitt.edu
Prof. Chenghua Lin	Long-term collaborator, The University of Manchester · chenghua.lin@manchester.ac.uk
Prof. Haoteng Tang	Collaborator, The University of Texas Rio Grande Valley · haoteng.tang@utrgv.edu