

Zaifu Zhan

zhan8023@umn.edu | [Personal Website](#) | [Linkedin](#) | [GitHub](#)

EDUCATION

University of Minnesota, Twin Cities	Ph.D	in	Electrical and Computer Engineering	GPA: 3.9/4.0	2021 – 2026
Tsinghua University	M.Eng.	in	Electrical Engineering	GPA: 3.6/4.0	2018 – 2021
Beijing Jiaotong University	BS	in	Electrical Engineering	GPA: 87/100	2014 – 2018

SKILLS

-
- **Programming:** Python(Proficient), MATLAB(Proficient), SQL, R, C/C++, Java
 - **Frameworks:** PyTorch, TensorFlow, Scikit-learn, Hugging Face
 - **Math:** Linear Algebra, Probability & Statistics, Convex and integer optimization, Common ML Algorithms, DL basics
 - **Data Processing:** Pandas, NumPy, Data Cleaning, Data Visualization (Matplotlib, Seaborn)
 - **Others:** Linux Bash, Shell, Git, GitHub, Slurm, LaTeX, Processing, Jupyter

RELATED EXPERIENCES [More experiences in [Linkedin](#)]

Graduate Research Assistant, UMN

Sep. 2021 – present

- EPEE: Entropy-based and patience-based early exiting for language models
 - Proposed the novel early exiting method that uses entropy and patience criterion to make early exit decision
 - Conducted comprehensive experiments on 6 biomedical datasets across BERT, ALBERT, and GPT2 models
 - Extended the method to the computer vision models ResNet18 and ResNet50 on three image datasets.
- Retrieval augmented multi-task learning LLMs for dietary supplements (DSs) [\[Paper\]](#)
 - Being the first work to do information extraction with LLMs for DSs via 4 tasks: name entity recognition, relation extraction, triple extraction, and usage status classification over 8 state-of-the-art LLMs such as Llama and Mistral.
 - Compared with single-task finetuning, multi-task finetuning LLMs achieve storage efficiency.
 - Applied MedCPT, Contriever, and BMRetriever to augment multi-task LLMs via a one-shot setting.
- Benchmarking Retrieval-Augmented LLMs in Biomedical NLP [\[Paper\]](#)
 - Fine-tuned Retrieval-Augmented Large Language Models (RALs) using LoRA across five biomedical tasks: triple extraction, link prediction, classification, question answering, and natural language inference, utilizing 9 biomedical datasets.
 - Benchmark RAL performance in key areas, including unlabeled robustness, counterfactual robustness, diverse robustness, and negative awareness.
 - Collaboratively reviewed state-of-the-art LLM research in disease diagnosis and co-authored a [\[paper\]](#).
- Optimized dataset combinations for multi-task learning via reinforcement learning
 - Generated random combinations from 12 datasets and fine-tuned the Llama3 model to collect combination-F1 score pairs, across four tasks: named entity recognition, relation extraction, event extraction, and classification.
 - Used a multi-layer neural network to predict the best combination, fine-tuned the LLM, and iteratively optimized dataset combinations.
- Computing stabilizing linear controllers via policy iteration [\[Github\]](#)
 - Proposed an iterative Q-learning reinforcement learning algorithm to find the optimal controller for noisy LTI systems using the action-value Bellman equation, demonstrating effectiveness under high noise conditions.
 - Transformed the core learning step into a least-squares problem, achieving high computational efficiency.
 - Proved the convergence of the Q-learning algorithm using conditional probability.
- Adversarial Learning Project: Out of Distribution Detectors vs. Attackers [\[Github\]](#)
 - Generated adversarial images from the CIFAR-10 dataset using over 30 attack methods.
 - Evaluated the performance of 18 detectors under all attacks using metrics like AUROC.
 - Automated the generation of shell scripts to run experiments on multiple GPUs for parallel computing.
 - Developed a novel two-sided threshold method, improving AUROC scores.

SELECTED PUBLICATIONS [\[Google scholar\]](#)

-
- [1] **Zhan, Z.**, & Zhang, R. (2025). Towards Better Multi-task Learning: A Framework for Optimizing Dataset Combinations in Large Language Models. In Findings of the Association for Computational Linguistics: NAACL 2025. (Accepted)
 - [2] **Zhan, Z.**, Zhou, S., Li, M., & Zhang, R. (2025). RAMIE: retrieval-augmented multi-task information extraction with large language models on dietary supplements. *Journal of the American Medical Informatics Association*, ocaf002.
 - [3] Li, M., **Zhan, Z.**, Yang, H., Xiao, Y., Huang, J., & Zhang, R. (2024). Benchmarking Retrieval-Augmented Large Language Models in Biomedical NLP: Application, Robustness, and Self-Awareness. *arXiv preprint arXiv:2405.08151*.