Zhizhang (Kevin) Hu

Contact Information

SE2 Building, 5200 Lake Rd
University of California, Merced
Merced, CA 95340 USA

E-mail: zhu42@ucmerced.edu
Homepage: kevin-hu.com
GitHub: github.com/KevinHooah/

Education

University of California, Merced, Merced, California USA **2020 - 2024** *(Expected Nov.)*

Ph.D. Candidate, Electrical Engineering and Computer Science
Research Interests:

- Multimodal Deep Learning, Causal Learning, Foundation Models for Science, Smart Healthcare

Academic Service:

- Reviewer: WACV 2024, ICCV 2023, IEEE Transactions on Parallel and Distributed Systems

Carnegie Mellon University, Pittsburgh, Pennsylvania USA **2020**

M.S., Building Science

- Thesis: Uncertainty Analysis of Electricity Load Prediction based on Bayesian Deep Learning

Selected Publications

Hu, Zhizhang, Xinliang Zhu, Son Tran, Rene Vidal, and Arnab Dhua. "ProVLA: Compositional Image Search with Progressive Vision-Language Alignment and Multimodal Fusion." In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshop, 2023.

Hu, Zhizhang, Shasha Li, Ming Du, Arnab Dhua, and Doug Gray. "Token Pruning-based Online Text Cleaning Guided by Multimodal Alignment for E-commerce Product Search." *In submission.*

Hu, Zhizhang*, Amirmohammad Radmehr*, Yue Zhang, Shijia Pan, and Phuc Nguyen. "IOTeeth: Intra-Oral Teeth Sensing System for Dental Occlusal Diseases Recognition." In Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp), 2024.

Hu, Zhizhang, Yue Zhang, Tong Yu, and Shijia Pan. "VMA: Domain Variance- and Modality-Aware Model Transfer for Fine-Grained Occupant Activity Recognition." In Proceedings of the ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), 2022.

Hu, Zhizhang, Tong Yu, Ruiyi Zhang, and Shijia Pan. "CIPhy: Causal Intervention with Physical Confounder from IoT Sensor Data for Robust Occupant Information Inference." In Proceedings of the ACM Conference on Embedded Networked Sensor Systems (SenSys) Workshop, 2022.

Professional Experience

Applied Scientist Intern
Amazon.com Inc. (Search Science and AI Group) **05/2023 - 08/2023**

- Work on large vision-language models for visual-grounded multimodal retrieval.
- Propose a token pruning-based algorithm for improving the multimodal embedding alignment given noisy image-text pairs.
- Implement a Ray-based distributed training pipeline for efficient large model training and experiment management.
- Improve the product visual search's recall rate up to 5 percentage points on the internal evaluation set, compared with deployed in-production models.

Applied Scientist Intern
Amazon.com Inc. (Visual Search & AR Team) **05/2022 - 08/2022**

- Propose a transformer-based vision-language fusion encoder for text-guided compositional multimodal retrieval.
- Introduce a momentum distillation-based hard negative mining strategy for a GPU RAM-efficient and data noise-robust deep metric learning.
- Improve the multimodal product search recall rate above the SOTA models on both external (Fashion 200K and Shoes) and internal product datasets.

Selected Academic Projects

Multimodal Learning for Ubiquitous Human Sensing
University of California, Merced **2020 -** *Present*

- Propose a multi-view multi-task deep learning model to embed and fuse the knowledge from wearable and infrastructure sensing modalities for fine-grained human activity recognition.
- Introduce a multimodal transfer learning framework that injects the physical knowledge to guide the model transfer for reducing the labeling cost and improving the model generalizability.
- Present a cross-granularity hierarchical learning algorithm that learns from the semantic relationship between data with different labeling quality and granularity for reducing the data annotation overhead while maintaining the model performance.

Foundation Models for Low-cost Air Quality Monitoring
University of California, Merced **2024 -** *Present*

- Benchmark large language models' (LLM) capabilities in assisting the decision-making of low-cost air quality sensors to improve large-scale PM 2.5 monitoring accuracy in poverty-strike areas.
- *On-going*: Exploring fine-tuning methods for improving LLM's multimodal capabilities on air quality analysis and decision-making assistance.

Data Properties Assessment and Intervention for Robust Inference
University of California, Merced **2020 - 2023**

- Investigate the impact of dataset inherent bias on the model generalization capability with causality-base fairness analysis.
- Propose a causal intervention method to mitigate the negative impact caused by confounding bias in the dataset on the person identification task.
- Introduce a framework to investigate the dataset transferability for assisting the transfer learning algorithm design. The framework conducts the analysis from distribution distance, information-dependency, and fairness perspectives, respectively.

Digital Healthcare for In-home Dental Occlusal Disease Monitoring
University of California, Merced; Yosemite Dental **2021 -** *Present*

- Developed an in-mouth sensing system to collect oral activities-induced teeth vibration data for long-term and low-cost dental occlusal disease monitoring.
- Introduce a physio-aware network block to augment the model's attention to channels that are indicative of the occlusal disease, which improves the model's generalizability to unseen users.

Honors and Awards

Best Poster Award for IPSN 2022	2022
Best Paper Award for the ACM UbiComp2020 Nurse Care Activity Recognition Challenge	2020
University of California Graduate Division NRT Fellowship	2020,2021
ACM International Workshop on Device-Free Human Sensing (DFHS) Travel Grant	2019
Carnegie Mellon University CFA Merit Scholarship	2018,2019

Technical Skills

Data science: PyTorch, TensorFlow, Transformers, TIMM, Ray, Sci-Kit Learn, SciPy, XGBoost, Prophet, LibROSA, NLTK, DASK, Bokeh, Holoviews

Programming Languages: Python, MATLAB, C, Shell, SQL, LaTeX