

# Zhizhang (Kevin) Hu

---

CONTACT INFORMATION	SE2 Building, 5200 Lake Rd University of California, Merced Merced, CA 95340 USA	<i>E-mail:</i> <a href="mailto:zhu42@ucmerced.edu">zhu42@ucmerced.edu</a> <i>Homepage:</i> <a href="http://kevin-hu.com">kevin-hu.com</a> <i>LinkedIn:</i> <a href="https://www.linkedin.com/in/zhizhang-hu/">linkedin.com/in/zhizhang-hu/</a>	<i>Google Scholar</i> <i>DBLP</i> <i>GitHub</i>
EDUCATION	<b>University of California, Merced</b> , Merced, California USA <b>2020 - 2024</b> ( <i>Expected Nov.</i> ) Ph.D. Candidate, Electrical Engineering and Computer Science Research Interests: <ul style="list-style-type: none"><li>• Multimodal Deep Learning, Causal Learning, Foundation Models for Science</li></ul> Academic Service: <ul style="list-style-type: none"><li>• Reviewer: WACV 2024, ICCV 2023, IEEE Transactions on Parallel and Distributed Systems</li></ul> <b>Carnegie Mellon University</b> , Pittsburgh, Pennsylvania USA <b>2018 - 2020</b> M.S., Building Science <ul style="list-style-type: none"><li>• Thesis: Uncertainty Analysis of Electricity Load Prediction based on Bayesian Deep Learning</li></ul> <b>Southwest Jiaotong University</b> , Chengdu, China <b>2014 - 2018</b> B.E., Mechanical Engineering		
SELECTED PUBLICATIONS	<b>Hu, Zhizhang</b> , Shasha Li, Ming Du, Arnab Dhua, and Doug Gray. “De-noised Vision-language Fusion Guided by Visual Cues for E-commerce Product Search.” In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshop, 2024. <b>Hu, Zhizhang</b> , 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. <b>Hu, Zhizhang*</b> , 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. <b>Hu, Zhizhang</b> , 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.		
PROFESSIONAL EXPERIENCE	<b>Amazon.com Inc.</b> Applied Scientist Intern ( <b>2022:</b> Visual Search & AR Team; <b>2023:</b> Search Science and AI Group) <ul style="list-style-type: none"><li>• Work on large vision-language foundation models for visual-grounded multimodal retrieval.</li><li>• Propose a token pruning-based algorithm for improving the multimodal embedding alignment given noisy image-text pairs.</li><li>• Propose a vision-language dual-attention encoder for compositional multimodal retrieval.</li><li>• Implement a Ray-based distributed training pipeline for efficient large-model training.</li><li>• Curate a foundation model fine-tuning dataset consisting of over 710,000 image-text pairs.</li><li>• Achieve the SOTA product visual search’s recall rate on the internal evaluation set ( compared with in-production models) and external (Fashion 200K and Shoes) datasets.</li><li>• Launched one model into real-world A/B testing.</li></ul>		
SELECTED ACADEMIC PROJECTS	<b>Multimodal Deep Learning for Ubiquitous Computing</b> <i>University of California, Merced</i> <b>2020 - 2023</b> <ul style="list-style-type: none"><li>• 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.</li><li>• Introduce a multimodal transfer learning framework that injects the physical knowledge to guide the model transfer to reduce the labeling cost and improve the model generalizability.</li><li>• Propose a causal intervention method to mitigate the negative impact caused by confounding bias in the dataset on the downstream inference tasks.</li></ul> <b>Large Language Models for Tabular and Sensing Data Reasoning</b> <i>University of California, Merced</i> <b>2024 - Present</b> <ul style="list-style-type: none"><li>• Propose the ” Order-of-Thought ” evaluation framework to evaluate the independence and guidance needed for large language models (LLMs) to conduct causal data reasoning.</li><li>• <i>On-going:</i> Propose a knowledge injection-based prompting framework to enhance LLMs’ capability in data reasoning for complex causal question answering.</li></ul>		
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, L <sup>A</sup> T <sub>E</sub> X		