

Xiwen (Christina) Wei

☎ +1 7348829746 ✉ xiwenwei@utexas.edu 🏠 Austin, TX 🌐 [Linkedin](#) 🌐 [Website](#)

Education

PhD in Electrical and Computer Engineering *University of Texas at Austin* **Austin, TX** Aug 2023-Dec 2027 (expect)

GPA: 3.8/4.00 (Cockrell School of Engineering PhD Fellowship)

BSE in Electrical Engineering *University of Michigan, Ann Arbor*

Ann Arbor, MI Aug 2021-May 2023

Summa Cum Laude, Dean's List, James B. Angell Scholar (GPA: 4.00/4.00)

BSE in Electrical and Computer Engineering *Shanghai Jiaotong University* **Shanghai, China** Sep 2019-Aug 2023

Outstanding Graduate of Shanghai Jiaotong University

Professional Experience

Graduate Research Assistant *System Level Design Group*

Austin, TX 08/2023 - present

- Working on **multimodal continual learning**, with a focus of preserving zero-shot generalization of **multimodal LLMs**.
- Designed an online task-free **continual learning** algorithm utilizing **efficient fine-tuning** of **foundation vision transformers**.
- Analyzed the fairness implications of **machine unlearning** in **diffusion models**. Developed a Bayesian optimization method to reduce model bias, balancing **privacy** preservation and **fairness**.

PhD Research Associate *Advanced Micro Devices, Inc. (AMD)*

Austin, TX 05/2025 - 12/2025

- Defined and led an end-to-end research project on fine-tuning of a **37B unified multimodal model**, targeting modality imbalance in interleaved text-image generation.
- Designed and implemented distributed supervised fine-tuning pipelines using **PEFT (LoRA)**, scaling efficiently across **8x MI300X GPUs** on AMD HPC clusters.
- Proposed a Pareto-based modality balancing algorithm, achieving up to **40%** improvement in image generation performance with minimal degradation in text generation quality.

Research Fellow, *University of Michigan Transportation Research Institute*

Ann Arbor, MI 01/2022 - 04/2023

- Developed a 3D parametric human model that represents diverse body types, enabling personalized and adaptive safety designs.
- Developed statistical models for thoracic spine geometry in MATLAB and R using Generalized Procrustes Analysis, Principal component analysis. Developed and analyzed **feedforward neural networks** to improve the predictive model.
- Processed medical images using **Mimics and HyperMesh** to quantify 3D geometries of human skeletons and internal organs.

Undergraduate Research Assistant, *Michigan Integrated Circuits Lab(MICL)* **Ann Arbor, MI** 05/2022 - 08/2022

- Designed a PID-based control algorithm in C and simulated the timer module in Michigan Micro Mote (M3) miniature sensor chip in MATLAB. Reduced timing error under extreme weather conditions by **83%**.
- Developed Python scripts (**PySerial, Pandas, Numpy**) for automated hardware verification.

Supply Chain Engineer Intern, *Soudronic AG*

Guangzhou, Guangdong, China 12/2020 - 05/2021

- Built a real-time inventory analytics system using Python (**Pandas, Scikit-learn**) and **SQL** to streamline order processing and enable proactive decision-making.

Publication

- Xiwen Wei**, Mustafa Munir, Radu Marculescu. *Mitigating Intra- and Inter-modal Forgetting in Continual Learning of Unified Multimodal Models*. Accepted by Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS), 2025.
- Xiwen Wei**, Guihong Li, Radu Marculescu. *Online-LoRA: Task-free Online Continual Learning via Low Rank Adaptation*. Accepted by IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 2025.
- Xiwen Wei**, Guihong Li, Radu Marculescu. *Fairness Implications of Machine Unlearning: Bias Risks in Removing NSFW Content from Text-to-Image Models*, NeurIPS 2024 Workshop on Regulatable ML.
- Yuedong Yang, **Xiwen Wei**, Mustafa Munir, Radu Marculescu. *Fuel Gauge: Estimating the Length of Chain-of-Thought a priori for Large Multi-modality Models*. Accepted by IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2026.
- Tianrui Hu, Dimitrios Liakopoulos, **Xiwen Wei**, Radu Marculescu, Neeraja J Yadwadkar. *Simulating rumor spreading in social networks using llm agents*, WMAC 2025: AAAI 2025 Workshop on Advancing LLM-Based Multi-Agent Collaboration.
- Mustafa Munir, Md Mostafijur Rahman, **Xiwen Wei**, Yuedong Yang, Radu Marculescu. *SearchViG: Optimal Vision GNNs via Ramanujan Spectral Optimization*. Accepted by The Fourth Learning on Graphs Conference (LOG), 2025.
- Mustafa Munir, **Xiwen Wei**, Harsh Goel, Minkyu Choi, Kartikeya Bhardwaj, Paul N. Whatmough, Sandeep P. Chinchali, Radu Marculescu. *ObjectAlign: Neuro-Symbolic Object Consistency Verification and Correction*, Under review.

Skills

- Machine Learning & AI:** Generative AI, Multimodal Large Language Models (MLLM), Continual Learning, Large Language Models (LLM), LLM-based agents, Time Series Forecasting, Federated Learning, Neural Networks.
- Deep Learning Frameworks:** PyTorch, TensorFlow, HuggingFace Transformers, PEFT, NumPy, Pandas, TFLite, Keras, Scikit-learn
- Programming Languages:** Python, C, C++, MATLAB, Assembly, Bash, SQL, SystemVerilog