Ziyu Zhao

(1)352-871-1917 ziyuz@email.sc.edu ziyuz-vision.github.io

SUMMARY

Applied Machine Learning Researcher (PhD) specializing in **Computer Vision** and **Multimodal AI**. Bridges research (**image processing, generative models**) with end-to-end production engineering from prototyping and model optimization (**LLM/ViT fine-tuning**) to cloud deployment (**AWS, Docker**). Adapts techniques across domains with focus on robustness, efficiency, and ethical impact. Solves real-world challenges where visual intelligence meets multimodal reasoning.

TECHNICAL SKILLS

- **LLM & Multimodal AI**:
  - **Fine-tuning**: Qwen (Qwen-VL, Qwen-1.8B), Llama 3, LoRA/QLoRA (Unsloth, Axolotl)
  - **RAG system**: LangChain, FAISS, cross-modal retrieval
  - **Models**: CLIP, BLIP-2, ViT, Stable Diffusion
- **Big Data & MLOps**: PySpark (ETL for multimodal datasets), AWS, Git, Docker, Linux, CI/CD, REST APIs, Unit Testing
- **Core Engineering**: Python (PyTorch, Pandas, NumPy), SQL, Bash, Linux, Unit Testing, Multi-GPU Training (DDP)

PROFESSIONAL EXPERIENCE

Reviewer : CVPR, ICCV, AAAI, ECCV, ACM MM, TMM, TPAMI, etc.

Research Scientist Intern in PAII Inc. *Jan 2025 – Aug 2025*

- ***LoRA-Fine-tuned Multimodal Generator for Brand-Compliant Visual Content***
  - Fine-tuned **Qwen-VL** with **rank-8 LoRA adapters** to generate brand-compliant marketing assets from structured prompts, reducing manual design effort by 70%
  - Engineered **image-conditioned diffusion pipeline** (Stable Diffusion 3.5 + ControlNet) to preserve brand identity across variations
  - Solved brand safety challenge: Integrated rule-based filters (via RAG) to block non-compliant generations (e.g., incorrect logo placement), achieving 99.2% policy adherence
- ***Multimodal RAG System for Compliance-Aware Content Generation***
  - Built RAG pipeline combining CLIP-based indexer for 100K+ compliant marketing assets and **Qwen-7B policy validator** fine-tuned on finance compliance rules (trained on 5K labeled examples from legal team)
  - Achieved **89% reduction in compliance violations** vs. vanilla diffusion models during A/B testing with 50K end-users

Graduate Instructional Assistant, University of South Carolina

- Led labs in **Algorithms**, **Data Structures**, and **Big Data Analytics** using **Python and Java**.
- Conducted **code reviews and debugging sessions** to help students optimize algorithmic efficiency.
- Supported courses in **Computer Architecture** and **Embedded Systems**.

RESEARCH EXPERIENCE

Dynamic Scene Reconstruction with 4D Gaussian Splatting *Mar 2025 – Aug 2025*

- **Developed a unified two-stream pipeline** integrating Hexplane-based scene modeling with SMPL-bound human avatars for high-fidelity reconstruction from monocular video.
- **Fused 6D Hexplane features with SMPL pose encodings** to drive a deformation decoder, enabling pose-controllable synthesis and temporal coherence.
- **Decoupled background and avatar dynamics** to optimize motion transitions and rendering fidelity, outperforming SOTA monocular methods in complex scene synthesis.

DPSeg: Dual-Prompt Cost Volume Learning for Open-Vocabulary Semantic Segmentation *Aug 2024 - Nov 2024*

- Built multimodal segmentation system integrating **vision-language models** for zero-shot object recognition, **directly transferable to product attribute extractio**n and **catalog enrichment** scenarios
- Designed **dual-prompt fusion mechanism** combining text and visual embeddings (+12% accuracy over baseline), applicable to **multimodal product search** where users describe items in natural language
- Optimized inference throughput by 30% through vectorized operations and mixed-precision training, critical for **real-time e-commerce applications** handling millions of queries

Cross-modal Few-shot 3D Point Cloud Semantic Segmentation via View Synthesis *Jan 2024 – May 2024*

- Firstly created cross-modal perception system combining **multi-view synthesis** and **vision-language embeddings** for few-shot learning for catalog expansion where new product categories need rapid onboarding with minimal examples
- Implemented **2D-to-3D knowledge transfer framework** reducing annotation requirements by 80%, directly applicable to **scaling product understanding systems** across diverse item categories

point cloud semantic segmentation based on class-specific Transformer network *May 2023 – Nov 2023*

- Designed **multi-scale transformer architecture** for efficient 3D feature aggregation reducing model parameters by **30%** while maintaining accuracy.
- Implemented **hierarchical attention modules** improving inference speed by 15% and enabling real-time processing of 100K+ point clouds.
- Developed **class-specific attention mechanism** reducing dependency on labeled data by **70%** for few-shot scenarios.

SELECTED PUBLICATIONS

- **Zhao, Z.**, Li, X., Shi, L., Imanpour, N., Wang, S.
"DPSeg: Dual-Prompt Cost Volume Learning for Open-Vocabulary Semantic Segmentation."
Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR), 2025, pp. 25346-25356

- **Zhao, Z.**, Li, X., Zhang, C., Cai, P., Wang, S.
"Crossmodal Few-shot 3D Point Cloud Semantic Segmentation via View Synthesis."
Proceedings of the 32nd ACM International Conference on Multimedia (ACM MM), 2024, pp. 2345–2353.

- Zhang, C., Wu, Z., Wu, X., **Zhao, Z.**, Wang, S.
"Few-Shot 3D Point Cloud Semantic Segmentation via Stratified Class-Specific Attention Based Transformer Network."
Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), vol. 37, no. 3, 2023, pp. 3410–3417.

- **Zhao, Z.**, Wu, Z., Wu, X., Zhang, C., Wang, S.
"Crossmodal Few-shot 3D Point Cloud Semantic Segmentation."
Proceedings of the 30th ACM International Conference on Multimedia (ACM MM), 2022, pp. 4760–4768.

- **Zhao, Z.**, Li, X., Zhang, C., Cai, P., Wang, S.
"Leveraging Adaptive Implicit Presentation Mapping for Ultra High-Resolution Image Segmentation." arXiv

EDUCATION

- **Ph.D. in Computer Science**, University of South Carolina *(Sep 2021 – **Expected** May 2026)*
- **M.S. in Computer Engineering (GPA: 3.72 / 4.0)**, University of Florida *(Aug 2019 – May 2021)*
- **B.S. in Mechanical Engineering (GPA: 91 / 100)**, Xidian University *(Sep 2014 – May 2018)*