Hanzhou Liu

PhD Student in **Computer Vision** | **LLM Post-Training & Tool-Augmented Reasoning**
heyhanzhou@gmail.com • 979 402 9122 • Google Scholar • LinkedIn • hanzhouliu.github.io

Technical Skills

LLM & VLM: Supervised Fine-Tuning, Post-Training, PPO, GRPO, RAG, Tool-Augmented Reasoning
Computer Vision & 3D: Transformers, SSMs, Diffusion Models, Image Restoration & Stylization, 3DGS
Programming & Frameworks: Python, PyTorch, Lightning, Hugging Face Transformers, VERL

Education

Texas A&M University, PhD in Computer Engineering — Aug 2021 – Expected 2027

Texas A&M University, MS Studies in Computer Engineering *(transferred into PhD)* — Aug 2019 – May 2021

Jilin University, BS in Electrical Engineering — Aug 2014 – June 2018

Work Experience

Research Assistant, Urban Resilience Lab, College Station — Aug 2025 – Jan 2026

- Developed ***FloodSQL-Bench***, the first Text-to-SQL benchmark for flood-risk analytics, comprising 443 question–SQL pairs across six difficulty levels and 10 tables with key-based, spatial, and hybrid joins.
- Designed a **metadata-driven, multi-granularity RAG pipeline** with table- and column-level retrieval, and evaluated 20+ proprietary and open-source LLMs on multi-table geospatial reasoning.

Applied Scientist Intern, Amazon Store, Seattle — May 2026 – Aug 2026

- Built the organization's first **tool-augmented reasoning system** for seller-enforcement decisions, using a **two-stage search process** to retrieve supporting evidence for decision verification and rationale generation.
- Developed an **LLM post-training pipeline** for seller-enforcement reasoning, combining synthetic reasoning traces, supervised fine-tuning, and **GRPO-based reinforcement learning** to improve rationale quality.

Selected Research Projects

Stylos: Multi-View 3D Stylization with Single-Forward Gaussian Splatting — Sept 2025 – Jan 2026

*Accepted by ICLR' 26 · **Top 1.3%** · **review scores (8/8/6/6)** | Paper Link | Code Link | Demo Link | Project Page*

- Developed ***Stylos***, a feed-forward 3D Gaussian Splatting framework that generates stylized 3D scenes from unposed single- or multi-view images without per-scene optimization or precomputed camera poses.
- Designed a dual-path Transformer architecture with **style–content cross-attention** and a novel **voxel-based 3D style loss** to preserve scene geometry and enforce view-consistent stylization across unseen scenes, and styles.

DeblurFlow: Progressive Image Deblurring with Video Diffusion Models — Jun 2025 – Aug 2025

Collaborative Research Exploration

- Constructed a training dataset of approximately **5K blurred–sharp image pairs** and converted them into progressive restoration videos using multiple temporal sampling schedules.
- Fine-tuned the ***Wan2.1*-14B image-to-video model** using LoRA to generate temporally coherent blur-to-sharp trajectories from a single blurred image.

XYScanNet: A State Space Model for Single Image Deblurring — Sep 2024 - Mar 2025

Accepted by CVPR-W' 25 | Paper Link | Code Link | Demo Link

- Proposed a novel **slice-and-scan *Vision State Space Module*** that addresses spatial misalignment in prior Mamba-based vision models by interleaving local intra-slice and global inter-slice scanning.
- Reduced **training time by 56.6%** and **inference memory usage by 46.2%** over the prior MambaIR module, while achieving a **17% lower KID** than the next-best deblurring method with competitive PSNR and SSIM.

Selected Publications

- **Stylos: Multi-View 3D Stylization with Single-Forward Gaussian Splatting**, ICLR 2026, top 1.3%, first author.
- **XYScanNet: A State Space Model for Single-Image Deblurring**, NTIRE at CVPR 2025, first author.
- **Mamba4Rec: Towards Efficient Sequential Recommendation with Selective State Space Models**, RelKD at KDD 2024, Best Paper Award, invited talk at Uber.