

Huy Huynh

Seattle, WA | (360)706-5778 | huy.cao.huynh.dev@gmail.com | LinkedIn | GitHub | Portfolio

EDUCATION

University of Washington, Seattle

B.S. in Computer Science

Expected Graduation: June 2027

GPA: 3.78

Coursework:

Deep Learning

Data Structures & Algorithms

Computer Vision

Systems Programming

Machine Learning

Linear Algebra

EXPERIENCE

Undergraduate Research Lead—Agentic Systems

January 2026 - Present

UW Systems Laboratory (SAMPL Group)

Advisors: Prof. Arvind Krishnamurthy | Mentor: Tapan Chugh

- Led development of an automated debugging and evaluation framework for autonomous web-agent systems, directing a four-person undergraduate research team from architecture to deployment.
- Architected an LLM-as-a-judge pipeline that classifies autonomous agent failures using retrieval-augmented prompting, structured error taxonomies, and human-annotated examples.
- Built a trace compilation pipeline that transforms raw JSON execution trajectories into token-efficient structured summaries for scalable LLM-based evaluation.
- Developed annotation schemas and evaluation workflows to analyze systematic failure modes in multi-step tool-using agents, accelerating research iteration and debugging.

Undergraduate Researcher | Generative Vision

June 2025 - Present

UW Graphics and Imaging Laboratory

Advisors: Profs. Steven Seitz, Brian Curless, Ira Kemelmacher | Mentor: Jingwei Ma

MicroZoom | First Author, SIGGRAPH Asia 2026

- Developed per-instance super-resolution methodologies, achieving state-of-the-art reconstruction at unprecedented scales (300×).
- Outperformed state-of-the-art baselines by 80.95% in image quality and 71.43% in reconstruction via a custom MultiDiffusion architecture.
- Preserved structural fidelity across giga-pixel generations by conditioning ControlNet adapters on low-resolution inputs and fine-tuned text tokens.
- Resolved texture ambiguity in low-resolution patches using SAM-extracted segmentation maps as distinct conditioning inputs.
- Cut trainable parameters by 99% and training time by 30% via LoRA adapter fine-tuning.

GarmentZoom | Third Author, ECCV 2026

- Achieved 74.6% top-1 user preference in visual quality and state-of-the-art perceptual scores (FID, LPIPS) across continuous 3-20× reference-guided super-resolution.
- Transferred high-fidelity textures from spatially unaligned references using VAE-encoded reference latents as a conditioning input.
- Enabled alignment-free texture learning by constructing a synthetic dataset of over 8,500 product images, using spatially separated crops as references to prevent trivial pixel-copying.

PROJECTS

LoRA-TTS: Efficient Dialect Adaptation | UW Advanced ML Course Project

December 2025

- Fine-tuned a transformer-based text-to-speech model for Taiwanese Mandarin using LoRA, reducing trainable parameters by 98%.
- Matched full fine-tuning synthesis quality at significantly lower GPU memory usage, trained on as little as 30 minutes of audio.
- Addressed data scarcity using transfer learning from a base Mandarin checkpoint, preserving prosody despite limited target-speaker data.

TECHNICAL SKILLS

Languages: Java, Python, C, C++, JavaScript, HTML/CSS

Libraries & Frameworks: PyTorch, React, NumPy, Pandas, OpenCV, scikit-learn, PIL, PostgreSQL

Tools: Git, GitHub, HuggingFace