

Jason Zhao

renjiz2@outlook.com | (206) 581-5662 | <https://renjiz2.com>

EDUCATION

B.S. in Computer Science (Data Science)

University of Washington Seattle Campus | Major GPA 3.80

Expected June 2027

Technical Skills:

Languages: Python, C/C++, CUDA, Java, C#, JavaScript, SQL, GoLang

ML / Systems: PyTorch, MuJoCo, Isaac Sim, ROS, NumPy

Areas: Computer Vision, Machine Learning, ML Systems, Reinforcement Learning, Distributed Systems

Research Publications:

Renjie Zhao, Jingwei Ma, Huy Huynh, Brian Curless, Steven M. Seitz, Ira Kemelmacher-Shlizerman.

“GarmentZoom: Generating Zoomable Images from Garment Listings”, ECCV 2026 RetailVision Workshop.

[\[Link\]](#) [\[arXiv\]](#)

Other Research:

Improving Physical Realism in Video Generation

February 2026 – Current

- UW RAIVN Lab – Advised by Prof. Ranjay Krishna
- Training Wan-2.2 5B with large scale reinforcement learning to improve physical realism in generated videos
- Developing scalable data-generation pipelines and automated verifiers

Power Adaptive Storage System

October 2024 – May 2025

- UW Systems Lab – Advised by Prof. Simon Peter
- Modified F2FS in Linux kernel to take advantage of the predictable power profiles of zone-namespaces storage devices
 - Stabilized IO performance with background garbage collection

Teaching:

Teaching Assistant — CSE 453: Data Center Systems

Fall 2026

- University of Washington · Instructor: Prof. Simon Peter

Projects:

ML-TOMB: Multi-LoRA Training with One-Model Base | Renjie Zhao, Alek Curless | [\[Report\]](#)

- Implemented a CUDA kernel for training multiple LoRA adapters on a shared base model
 - Achieved up to 2.63× training speedup

SiriTok - Semantically Rich 1D Tokenizer | Renjie Zhao | [\[Poster\]](#)

- A one-dimensional image tokenizer for latent diffusion models while preserving both textural details and semantic information
 - Achieved 79% linear probing accuracy on ImageNet
- Fine-tuned a pretrained semantic encoding model (DINOv3) with a semantic preservation loss