

BENRAN HU

bhuai@connect.ust.hk | zymk9.github.io | +1 (412) 224-7539 | San Jose, CA

EDUCATION

Carnegie Mellon University

Master of Science in Computer Science

Dec 2024
Pittsburgh, PA

The Hong Kong University of Science and Technology

Bachelor of Science in Data Science and Technology, and in Computer Science

Jun 2023
Hong Kong SAR

SKILLS

Languages: Python, C/C++, CUDA

GenAI/ML: diffusion models, transformers, tokenizer design, image/video generation, large-scale model training, data curation and recaptioning

3D Vision/Graphics: NeRF, 3D Gaussian splatting, differentiable rendering, 3D reconstruction, path tracing

Systems/Tools: PyTorch, CUDA, Kubernetes, Slurm, Hugging Face (Diffusers, Transformers, Accelerate), Linux, Git

EXPERIENCE

Applied Research Engineer

GenAI Group, AMD

Mar 2025 – Present
San Jose, CA

- Develop text-to-image generative models built on high-compression 1D diffusion tokenizers, spanning model development, training, data preparation, and release.
- Design and build large-scale data curation and recaptioning pipelines for high-quality text-to-image training datasets.
- Support the development of unified vision-language models for joint visual understanding and image generation.

Research Intern

Creative Vision Group, Snap Research

Jun 2024 – Aug 2024
Santa Monica, CA

- Improved transformer-based video autoencoders and latent diffusion models for video and image generation.
- Boosted post-training performance of a large-scale video generative model through checkpoint aggregation.

RESEARCH EXPERIENCE

3D Reconstruction with Improved Dipole Sums

CMU, Present

- Improve regularized dipole sums for faster optimization, robustness to sparse initialization, and higher-quality 3D surface reconstruction from multi-view inputs.

Uncertainty Quantification in Differentiable Rendering ^[web]

CMU, 2024

- Developed a general and efficient framework to quantify aleatoric and epistemic uncertainty in both parameter and image space for inverse rendering tasks via uncertainty propagation.

Segment Anything for NeRF [CVPR'24] ^[pdf] ^[web]

HKUST, 2023

- Combined Segment Anything and Neural Radiance Fields to enable high-quality interactive 3D object segmentation.
- Explored distillation strategies for vision foundation models to improve NeRF segmentation quality.

NeRF 3D Instance Segmentation [ICCV'23] ^[pdf] ^[code]

HKUST, 2023

- Introduced one of the first 3D instance segmentation methods in NeRFs by optimizing a neural field of instance masks.
- Built an end-to-end 3D segmentation pipeline on voxel representations based on 3D Mask R-CNN.

3D Object Detection in NeRF [CVPR'23] ^[pdf] ^[code]

HKUST, 2022

- Proposed NeRF-RPN, the first 3D object detection framework operating directly on Neural Radiance Fields from multi-view RGB inputs alone.
- Created the first large-scale synthetic indoor dataset for NeRF-based 3D object detection using Blender.

PUBLICATIONS

*equal contribution

Instella-T2I: Pushing the Limits of 1D Discrete Latent Space Image Generation | *arXiv*, 2025 [pdf]

Ze Wang, Hao Chen, [Benran Hu](#), Jiang Liu, Ximeng Sun, Jialian Wu, Yusheng Su, Xiaodong Yu, Emad Barsoum, and Zicheng Liu.

Improving the Diffusability of Autoencoders | *ICML*, 2025 [pdf]

Ivan Skorokhodov, Sharath Girish, [Benran Hu](#), Willi Menapace, Yanyu Li, Rameen Abdal, Sergey Tulyakov, and Aliaksandr Siarohin.

SANeRF-HQ: Segment Anything for NeRF in High Quality | *CVPR*, 2024 [pdf]

Yichen Liu, [Benran Hu](#), Chi-Keung Tang, and Yu-Wing Tai.

Instance Neural Radiance Field | *ICCV*, 2023 [pdf]

Yichen Liu*, [Benran Hu](#)*, Junkai Huang*, Yu-Wing Tai, and Chi-Keung Tang.

NeRF-RPN: A General Framework for Object Detection in NeRFs | *CVPR*, 2023 [pdf]

[Benran Hu](#)*, Junkai Huang*, Yichen Liu*, Yu-Wing Tai, and Chi-Keung Tang.

SELECTED PROJECTS

Point-based 3D Gaussian Manipulation with Diffusion [web]

3DGS, Diffusion Models, Scene Editing

- Developed a 3D scene editing pipeline combining point-guided image editing with diffusion models and 3D Gaussian splatting.

Wavefront Path Tracing [code]

CUDA, GPU Programming, Path Tracing

- Implemented a CUDA wavefront path tracer with wide BVH traversal, achieving 1.79x speedup over megakernel and 8.20x over multi-threaded CPU baselines.

Specular Manifold Sampling

Path Tracing, Caustics Rendering

- Integrated single-scattering Specular Manifold Sampling into a custom path tracer for efficient caustics rendering on specular-diffuse-specular light paths.

Vulkan Mesh Shading and Culling

Vulkan, Mesh Shading

- Built a high-performance meshlet shading and culling pipeline in Vulkan leveraging mesh shaders and task shaders.