

Yixing Lu

✉ yixing@cs.unc.edu | 🌐 lewis-lv0.github.io | 🎓 Google Scholar | 🐙 github.com/Lewis-lv0

SUMMARY

Ph.D. student at UNC Chapel Hill working on 3D/4D vision and generative models, with a current focus on human-object interaction for humanoid robots. Three ICCV/ECCV papers on 3D human reconstruction and generation, including a first-author ICCV 2025 paper. Industry research at Baidu ERNIE on large-scale pretraining of unified visual tokenizers for multimodal LLMs.

EDUCATION

University of North Carolina at Chapel Hill

Ph.D. in Computer Science, Advisor: Zhongzheng (Jason) Ren

Aug. 2026 – Present

Chapel Hill, USA

Carnegie Mellon University

M.S. in ECE, Advisor: Fernando De la Torre

Aug. 2023 – May 2025

Pittsburgh, USA

University of Liverpool

B.S. in Computer Science, First Class Honours

Sept. 2019 – Jun. 2023

Liverpool, UK

RESEARCH EXPERIENCE

RAIR Lab, UNC Chapel Hill

Ph.D. Researcher

Chapel Hill, USA

Aug. 2026 – Present

- Developing a synthetic data pipeline that combines conditional video generation with constrained 4D reconstruction to synthesize human interactions with articulated objects.
- Exploring cross-embodiment retargeting of synthesized human motions to humanoid robots.

ERNIE Pretraining Team, Baidu

Research Intern

Beijing, China

Feb. 2026 – Aug. 2026

- Pretrained visual tokenizers for unified multimodal models, replacing separate understanding and generation encoders with a single latent space via joint image-text contrastive, self-distillation, and reconstruction objectives.
- Scaled tokenizer pretraining to 1B image-text pairs on 128 GPUs, establishing that downstream generation quality scales with tokenizer pretraining compute, whereas reconstruction-only autoencoders saturate early.

Human Sensing Lab, CMU

Graduate Researcher, Advisor: Fernando De la Torre

Pittsburgh, USA

Sept. 2023 – May 2025

- Proposed a video diffusion-based method for single-image human synthesis with 3D representation guidance.
- Built a generalizable 3D Gaussian framework for sparse-view human reconstruction with geometric priors.
- Investigated native 3D generation of human Gaussians, scaling training to millions of reconstructed 3D assets.

PUBLICATIONS

* Equal contribution

1. **Yixing Lu**, Junting Dong, Youngjoong Kwon, Qin Zhao, Bo Dai, Fernando De la Torre. *GAS: Generative Avatar Synthesis from a Single Image*. **ICCV 2025**. [Paper](#)
2. Yuhang Yang*, Fengqi Liu*, **Yixing Lu**, Qin Zhao, Pingyu Wu, Wei Zhai, Ran Yi, Yang Cao, Lizhuang Ma, Zheng-Jun Zha, Junting Dong. *SIGMAN: Scaling 3D Human Gaussian Generation with Millions of Assets*. **ICCV 2025**. [Paper](#)
3. Youngjoong Kwon, **Yixing Lu***, Baole Fang*, Haoye Dong, Cheng Zhang, Francisco Vicente Carrasco, Albert Mosella-Montoro, Jianjin Xu, Shingo Takagi, Daeil Kim, Aayush Prakash, Fernando De la Torre. *Generalizable Human Gaussians for Sparse View Synthesis*. **ECCV 2024**. [Paper](#)
4. **Yixing Lu**, Zhaoxin Fan, Min Xu. *Multi-dimensional Fusion and Consistency for Semi-supervised Medical Image Segmentation*. **MMM 2024**. [Paper](#)