

Kuan Heng (Jordan) Lin

✉ kuanhenglin@ucla.edu

📞 (424) 440-5778

🌐 kuanhenglin

🌐 kuanhenglin.github.io

Skills

Research: Generative AI, computer vision, multimodal/3D generation and representation, model parallelism & scaling, computational imaging
Programming: Python (PyTorch, TensorFlow, JAX), CUDA, C++, C, Assembly, Verilog, R, Haskell, React, Flask, JavaScript, HTML/CSS

Education

Computer Science B.S., minor in Mathematics

September 2021–Present

University of California, Los Angeles (GPA: 3.981, Dean's Honor List, Upsilon Pi Epsilon)

Coursework: Computer Vision, Deep Learning, Graphics, Algorithms & Data Structures, Imaging, Signals & Systems, Programming Languages, Software Construction, OS, Quantum Computing, Systems, Architecture, Theory, Digital Design, Linear Algebra, Analysis, Probability, Statistics

Awards & Honors: 2024–25 [CRA Outstanding Undergraduate Researcher Award](#) Finalist, 2024–25 [Undergraduate Research Scholars Program](#) (\$6000), [NeurIPS 2024 Scholar Award](#), 2023 [URC-Sciences Summer Program](#) (\$6000)

Research

Publications | Generative modeling, computer vision, controllable generation, deep learning

December 2023–Present

[1] [Kuan Heng Lin*](#), Sicheng Mo*, Ben Klingher, Fangzhou Mu, Bolei Zhou. “[Ctrl-X: Controlling Structure and Appearance for Text-To-Image Generation Without Guidance](#)”. In: *Neural Information Processing Systems (NeurIPS)*. 2024. [[NeurIPS Scholar Award](#)]

[2] Sicheng Mo*, Fangzhou Mu*, [Kuan Heng Lin](#), Yanli Liu, Bochen Guan, Yin Li, and Bolei Zhou. “[FreeControl: Training-Free Spatial Control of Any Text-to-Image Diffusion Model with Any Condition](#)”. In: *Computer Vision and Pattern Recognition (CVPR)*. 2024.

Undergraduate Researcher | [Zhou Lab at UCLA](#)

March 2023–Present

- Design fast [training-free and guidance-free structure and appearance control](#) for arbitrary text-to-image and text-to-video models.
- Develop [training-free controllable generation](#) methods via guidance and feature injection for Stable Diffusion (SD) and video diffusion models for image and video translation and stylization, building large-scale PyTorch + Diffusers pipelines and benchmarks.
- Explore [semantic latent space manipulation of diffusion models](#), leveraging inversion to perform image editing latent guidance.

URC-Sciences Summer Program Scholarship Researcher | [Zhou Lab at UCLA](#)

June 2023–September 2023

- Design a [human-in-the-loop video generator](#) by extending Stable Diffusion with video guidance and interactive grounded generation.

Undergraduate Researcher | The Ozcan Research Group (HHMI Program)

October 2022–June 2023

- Design Fourier residual and attention blocks for diffusion autoencoders and generative adversarial networks for accurate and disentangled hologram reconstruction, super-resolution, and axial distance prediction at reduced network sizes.

Work Experiences

Research Intern | [Snap Inc.](#)

June 2024–September 2024

- Work in the [Creative Vision](#) team to optimize training and inference of large-scale video generation with PyTorch profiling and parallelism.
- Implement pipeline-wide variable sequence length training for arbitrary datasets and conditioning for flexible, efficient multimodal training.
- Build Snap's first distributed 2D parallel (data + model parallel) training framework for scaling video Diffusion Transformers with `torch.distributed`, designing new tensor and sequence parallel strategies to minimize VRAM usage and GPU communication overhead.

Program Development Team ← **Head LA** | [UCLA CS](#), [UCLA CAE](#), [Learning Assistant Program](#)

March 2023–Present

- Optimize LA application review with [Airtable](#) and [Gmail JavaScript](#) for 900+ applicants supporting 14000+ students in STEM courses.
- Advertise the LA program to increase applicants for CS courses by 300% and work with professors & TAs to implement evidence-based pedagogy and collaborative, inclusive learning in CS classrooms for high-demand upper-division courses (e.g., COM SCI 111, 118, 131).
- Lead weekly discussions and workshops for COM SCI 33: Computer Organization for 400+ students. Reviewed very positively.

Projects & Experiences

Advisor ← **Co-President** ← **Workshops Officer** | [ACM Student Chapter at UCLA](#), AI Committee

May 2022–Present

- Founded weekly [reading groups & seminars](#) discussing recent ML advances such as generative vision, reinforcement learning, and LLMs with student & industry speakers, culminating in projects and events for general members such as the adversarial AI competition.
- Spearhead general member programs (e.g., special topics discussion) and bold initiatives (e.g., research team, shared compute, inter-committee collabs, AI×Cyber Symposiums) which drastically improved member retention, officer burnout, and club exposure.
- Revamp and teach workshops on deep learning topics such as gradient descent, neural networks, CNNs, Transformers and generative vision.

Hackathons | PyTorch, TensorBoard, JavaScript, React.js, Flask, Solidity, Web3.js

January 2022–Present

- LA Hacks 2023 (Overall Third Place), [people2vec](#): Social media platform powered by LLMs, CV, and YouTube watch histories that matches people by their media interests. Integrated sentence embeddings and Inception V3 features for distribution analysis to compute interest similarity. Visualize user data with PCA “video cloud” to convey matched interests while preserving privacy.
- HackMIT 2022 (Blockchain for Society Second Place), [Wikisafe](#): Crowd-sourced knowledge database powered by machine learning and blockchain for secure version management. Integrated fine-tuned text summarization, caption generation, and generative imagery PyTorch models and Solidity smart contracts on the Ethereum blockchain with Web3.js in a full-stack web application.