Malcolm Hsiu

510-371-1758 | mhsiu@ucsd.edu | github.com/mhsiu01

Education

UC San Diego — La Jolla, CA

Bachelor of Science in Mathematics-Computer Science, ***GPA 3.725*** — *Oct. 2020 – June 2024*

Experience

Student Researcher - Computer Vision+Robotics — March 2024 – August 2024

Hao Su Lab @ UC San Diego — *La Jolla, CA*

- Researched methods for object **metric scale estimation** from monocular RGB data, for application in scene generation. Trained convolutional nets on a filtered subset of **Amazon Berkeley Objects** dataset with size metadata.
- Worked on **domain randomization** for the **ManiSkill** robotics simulator in Prof. Hao Su's lab. Added feature enabling objects to have randomized color and texture, and enabling lighting to have randomized color, brightness, and pose.
- Performed literature survey on interpretability of robotics models, especially bias in common visual encoders.

Software Engineer - Computer Vision — June 2022 – Sep. 2022

Accel Robotics — *San Diego, CA*

- Implemented and trained **YOLOv3 object detection** model to support customer tracker (for automated grocery/convenience store checkout) by recovering customers dropped by the primary customer tracker.
- Cleaned and labeled image data from store's overhead cameras.

Student Researcher - NLP — Oct. 2021 – June 2022

Shang Data Lab @ UC San Diego — *La Jolla, CA*

- Studied **document classification** with partially unknown classes under Prof. Jingbo Shang
- Clustered BERT document embeddings with **Gaussian mixture models**, visualized embedding space with GPU-accelerated **t-SNE** via RAPIDS library, and created simple frequency-based algorithm for recovering unknown classes based on known classes. Presented results in poster session to UCSD students and faculty.

Projects & Paper Implementations

Stable Diffusion | *Huggingface diffusers, PyTorch* — Oct. 2024 – Nov. 2024

- Used pre-trained **ControlNets** (canny-edge, QR code), **LORA**, and **Stable Diffusion 1.5** to create stylized logos. Implemented **parameter-efficient fine-tuning** on Stable Diffusion models: textual inversion and LORA.

Tensor Radiance Fields | *PyTorch, numpy, NeRF* — Oct. 2023 – Dec. 2023

- Implemented TensoRF (from *TensoRF: Tensorial Radiance Fields*), a **neural radiance field** variant that uses low-rank tensor factorization to efficiently approximate a dense voxel space. Implemented simple ray-marching, AABB collision detection.

Transformer Decoder | *einops, PyTorch* — June 2023 – Sep. 2023

- Transformer decoder from scratch, concisely implemented via **einsum/einops**
- Implemented multi-head self-attention, causal masking, char-level tokenizer, and learned token embeddings

Generative Modeling from Scratch | *PyTorch* — June 2023 – Sep. 2023

- Small UNet trained on **DDPM** diffusion objective to unconditionally generate MNIST numbers
- Simple **variational autoencoder** using small MLP, trained to unconditionally generate MNIST numbers, with visualization of latent space

Style Transfer from Scratch | *PyTorch* — June 2023 – Sep. 2023

- Implemented Gatys style transfer (from *A Neural Algorithm of Artistic Style*), which imbues a base image with a reference image's visual style using pre-trained convolutional net features.

Loss Landscape Visualization | *PyTorch, matplotlib* — June 2023 – Sep. 2023

- Visualizing effect of residual connections on the loss landscape of ResNets, as proposed in *Visualizing the Loss Landscape of Neural Nets*

Technical Skills

Languages: Python (advanced), Java, C++, PostgreSQL
Developer Tools: PyTorch, Distributed Data Parallel (DDP), Jupyter, Git, Docker, Kubernetes, LaTeX
Libraries: NumPy, Matplotlib, pandas, cv2, PIL, Huggingface diffusers, Dask, Tensorboard

Coursework

Machine Learning: Deep Learning for 3D Data (graduate course), Unsupervised Learning (graduate course), Probabilistic Search and Reasoning (graduate course), Deep Learning, Artificial Intelligence
Computer Science: Data Structures, Algorithms, Databases, Theory of Computation, Computer Organization
Math: Honors Linear Algebra, Honors Multivariable Calculus, Intro to Statistics, Intro to Probability, Random Matrix Theory, Linear Programming, Real Analysis, Abstract Algebra, Graph Theory and Extremal Combinatorics

Teaching / Outreach / Events

ACM AI - Events Board Member June 2023 – June 2024
UC San Diego *La Jolla, CA*

- Organized AI workshops and socials for UC San Diego's ACM chapter. Created presentations and accompanying notebook implementations to offer hands-on experience. Presented to audiences of several dozen students.
- Workshop topics: basics of machine learning, convolutional neural nets, style transfer, diffusion models, data poisoning for text2img models
- Implemented Gatys style transfer, which imbues a base image with a reference image's visual style using pre-trained convolutional net features.

Science Initiative - Internal Vice President Oct 2020 – June 2021
UC San Diego *Remote*

- Participated in STEM outreach for San Diego K12 students during COVID quarantine. Heavily involved in ideation, planning, and execution of various STEM-themed events held over Zoom.
- Created and demonstrated a variety of live chemistry and physics experiments for audiences of several dozen students. Helped organize and execute a science-themed *Jeopardy!* night.
- Recruited new org members using social media.

UC Berkeley AI Hackathon June 2024
UC Berkeley *Berkeley, CA*

- Created a self-narrating interactive history textbook.
- Used Hume AI's voice models to create realistic narrators grounded in any historical event the user specifies.
- I integrated Stable Diffusion to generate illustrations accompanying the narration.