

Kaustav Mukherjee

415-216-5804 | kaustav.mu@gmail.com | linkedin.com/in/kaustavmu | github.com/kaustavmu | kaustavmukherjee.com

EDUCATION

Carnegie Mellon University (CMU)

Master of Science in Computer Vision

Pittsburgh, PA

Aug. 2024 – Dec. 2025

National University of Singapore (NUS)

Bachelor of Engineering in Mechanical Engineering

Singapore, Singapore

Aug. 2020 – May 2024

EXPERIENCE

AI Engineer

Jan 2026 – Present

Kaliber Labs

San Francisco, CA

- Built a real-time 3D multi-view reconstruction pipeline from arthroscopic surgical video using TSDF fusion, directly resulting in a flagship product demo and a co-authored paper accepted at ECCV 2026 Workshop DCA-MI
- Developed novel architectures for JEPA-style world models that improve baselines by over 50% on 3D benchmarks
- Improved textual and aural emotion recognition models by architecting a multi-task learning (MTL) framework for categorical classification and regression, driving a +30% absolute jump in accuracy
- Engineered a stateless microservice architecture with NATS and Redis for a multimodal conversational AI PoC

AI Engineer Intern

May 2025 – Aug. 2025

Kaliber Labs

San Francisco, CA

- Architected a modular robotic navigation system with ROS2 and Nav 2, integrating 6+ CV models with C++ behavior trees via shared memory and multithreading, and tested on the LG Q9 robot and NVIDIA Isaac Sim
- Reduced per-frame edge processing latency from 300ms to 150ms across multiple CV inference tasks
- Led technical integration across 10+ team members and delivered 10+ live demos to investors and stakeholders

Software Engineering Intern

Jan. 2023 – Jul. 2024

Kaliber Labs

San Francisco, CA

- Accelerated convergence of a surgical camera pose estimation algorithm by 10 times using particle swarm optimization (PSO) to match semantic segmentation masks between the scope view and a digital twin.
- Developed an intra-operative camera registration algorithm using Neural Radiance Fields (NeRFs), enabling linear and geodesic measurements between 3D points with above 90% accuracy
- Fine-tuned stable diffusion to generate synthetic training data, expanding surgical tool datasets by up to 5 times
- Employed artificial data augmentation to train multiple semantic segmentation models to above 97% accuracy
- Optimized CV model inference pipelines via ONNX and quantization, reducing memory footprint by 40%

PUBLICATIONS

Bychkov, D., Shah, R., Mukherjee, K., So, M. "Surface Coverage Completeness for Quality-Assured Arthroscopic 3D Reconstruction" *Proceedings of the ECCV Workshop on DCA-MI, 2026 (To appear).*

PROJECTS

Bundle Adjustment for 360° Pose Correction | *Gaussian Splatting, SfM, SLAM*

Aug 2025 – Nov 2025

- Implemented a rig-based Structure from Motion (SfM) pipeline with bundle adjustment for Stella-VSLAM outputs from 360° images, improving Gaussian Splatting reconstruction quality by 10.5%.

AvatarsFTW: 3D Human Avatars From the Wild | *PyTorch, Image Generation*

Oct 2024 – Dec 2024

- Created pipeline to improve performance of SOTA 3D human avatar generation models on in-the-wild images
- Developed novel method for human keypoint estimation to condition diffusion models for human image generation

Drone Perception System at CMU Airlab | *Python, OpenCV, ROS, Computer Vision*

Aug 2024 – Nov 2024

- Designed a depth-aware perception system for a drone with a 2-DOF manipulator using ROS and RGBD cameras

SKILLS

Competencies: 3D Reconstruction, Computer Vision, SLAM, SfM, NeRF, Gaussian Splatting, Diffusion, World Models, Reinforcement Learning, LLM Fine-tuning, Context Engineering, Model Quantization

Frameworks: PyTorch, Tensorflow, NumPy, OpenCV, ROS2

Tools: Docker, Redis, NATS, ONNX, AWS, Linux, Git

Languages: Python, C++, JavaScript