

# OJAS TASKAR

ojastaskar22@gmail.com | +14434184243 | <https://github.com/ojas200> | [ojas200.github.io](https://ojas200.github.io)

## EDUCATION

---

JOHNS HOPKINS UNIVERSITY | Maryland, United States

**Master of Science, Robotics** | Expected 05/2026 | 3.8/4.0

**Coursework:** Computer Vision, Generative Models for Biomedical Vision, Machine Perception, Vision as Bayesian Inference

**Teaching Experience: TA:** Introduction to Robot Learning, Algorithms for Sensor Based Robotics, AI Essentials for Business

VEERMATA JIJABAI TECHNOLOGICAL INSTITUTE | Mumbai, India

**Bachelor of Technology, ECE** | 06/2024 | **8.78/10**

Bachelor of Technology in Electronics and Telecommunication with *Distinction*

## TECHNICAL SKILLS

---

**Languages:** Python, C, C++, Julia, CUDA, JAX, MATLAB, Git, Bash, JavaScript, HTML/CSS

**Robotics Packages:** CMake, GTSAM, ROS2, Gazebo, MoveIt, OMPL, Docker, MLflow

**Deep Learning:** PyTorch, TensorFlow, TensorRT, OpenCV, ONNX, HuggingFace, Slurm, AWS (EC2), CI/CD, Jetson Nano

**Simulators:** MuJoCo, NVIDIA IsaacSim, CoppeliaSim, Microsoft AirSim

## RESEARCH EXPERIENCE

---

**Research Intern** | Image Analysis & Communications Laboratory, JHU

*June 2025 - August 2025*

- Built a bidirectional diffusion model to synthesize paramagnetic rim lesions on healthy brain susceptibility-weighted images.
- Modified a SOTA multi-view DDIM inversion & fusion pipeline conditioned on lesion masks for 9.85x faster synthesis output.
- Performed cross-modality fusion of axial and coronal frames to ensure consistency of generated lesions across DDIM outputs.

**Graduate Research Assistant** | Biomechanical Image Guided Surgical Systems Lab, JHU

*September 2024 - June 2025*

- Developed multi-camera 6DoF pose estimation for sub-millimeter tracking of patient's head, for robot-assisted surgery.
- Built a Bayesian plane fitting algorithm to solve monocular image depth ambiguity, reducing depth error from 40 mm to 6 mm.
- Deployed deformable Gaussian Mixtures as PyTorch plugins for 3D point cloud to MRI registration, improving quality over ICP.
- Published a [manuscript](#) at IEEE VR 2026 on surgical instrument tracking on Microsoft HoloLens2 using pose graph optimization.

**Research Intern** | KIREAP, Inc. India

*January 2024 - June 2024*

- Prototyped algorithm development of a drone vision-based landing system, reducing AprilTag detection false positives by 20%.
- Utilized AprilTag pattern arrangement and lightweight CNNs to design passive pad tracking system, with 77% trial success.
- Benchmarked segmentation pipelines on Jetson Nano, writing custom CUDA kernels for 35% faster per frame recognition.

## PROJECTS

---

**Fine-tuning Flow Matching for Preference-Guided Generation** | *RL, Generative Models, Optimization, Model Deployment* [\[GitHub\]](#)

- Built a REINFORCE policy gradient algorithm to fine-tune pre-trained flow matching policy to align with user-defined prompt.
- Incorporated CLIP VIT-B/32 encoding to provide rewards using cosine similarity of prompt and generation in latent space.
- Evaluated fine-tuning quality and flow matching exploration by validating on CelebA dataset, achieving mean pSNR of 27.5 dB.

**Neuroimage Registration and Synthesis** | *3D Synthesis, Medical Image Analysis*

[\[GitHub\]](#)

- Formed a pipeline for diffeomorphic registration between T1w/T2w and FA/ADC diffusion tensor MRI, improving robustness to imaging artifacts and inter-modality resolution differences by post-training optimization on holdout registered pairs.
- Fine-tuned a pretrained SegResNet segmentation network to synthesize 3D volumes, achieving validation MSE of 0.0026.
- Curated a small dataset of skull-stripped brain MRIs to train a UNet for automated skull-stripping, with DICE score of 0.76.

**Contrastive Pre-Training for Object Segmentation** | *Semantic Segmentation, Self-supervised Learning, Pruning*

[\[GitHub\]](#)

- Wrote a [SimCLR](#)-based pre-training pipeline with  $\gamma$ -aware & InfoNCE loss for anomaly detection, with 42% validation accuracy.
- Developed Vision Transformer and Decision Tree-based classifiers with LoRa fine-tuning on pre-trained backbone, evaluating imbalanced classification performance using precision-recall curves on Tensorboard using PyTorch.
- Performed comparative analyses of data augmentation techniques on pre-training efficiency and impact on validation loss, determining gaussian blur, and random image crops having 78% correlation to positive detection.

**3D Reconstruction using Gaussian Splatting and NeRFs** | *3D Rendering, Real-time Inference, Model Profiling*

[\[Report\]](#)

- Trained classical NeRF, Instant-NGP NeRF, TensorRF and classical Gaussian Splatting for 3D reconstruction from video frames.
- Profiled GPU kernels with NVIDIA Nsight Compute and PyTorch Profiler, demonstrating real-time rendering gains via differentiable Gaussian rasterization, deployed using CUDA/C++ on an NVIDIA RTX3060 GPU.
- Evaluated fidelity of multi-view synthesis using PSNR and SSIM metrics, achieving 5.2x better clarity for Gaussian Splatting.