

# Mayesh Mohapatra

(979) 574-0276 | [mayesh212@gmail.com](mailto:mayesh212@gmail.com) | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)

## EDUCATION

### Texas A&M University

Master of Computer Science

### Bennett University

B.Tech (Honors), Computer Science Engineering — GPA: 3.54/4.0

College Station, TX

Aug. 2025 – May 2027 (Expected)

Greater Noida, UP

May 2023

## TECHNICAL SKILLS

**Languages:** Python, C/C++, SQL, R    **Frameworks:** PyTorch, TensorFlow, Scikit-Learn, NumPy, Pandas

**Perception & CV:** 3D Reconstruction, 4D Gaussian Splatting, Pose Estimation, Object Detection & Tracking, Segmentation, Human3R, VGGT, DeepLabv3, DINOv2, SMPL

**Tools & Infra:** CUDA, Docker, AWS, Spark, Weights & Biases, GDAL, Git, COLMAP    **Evaluation:** PSNR, SSIM, LPIPS, mAP, RMSE, AUC-ROC

## EXPERIENCE

### Graduate Researcher | 3D Perception & Scene Understanding

Aug. 2025 – Present

Texas A&M University | Advisor: Prof. Cheng Zhang

College Station, TX

- Developed an end-to-end 3D scene reconstruction pipeline from monocular athlete video, integrating **Human3R**, **VGGT** (pose estimation), and **DeepLabv3** (segmentation) into a unified tracking framework.
- Accelerated inference via mixed-precision CUDA kernels and INT8 quantization, achieving **40% latency reduction** and enabling real-time biomechanical feedback during live training.
- Engineered a module to auto-extract metrics (jump height, velocity, takeoff angle) for **data-driven coaching** across **3+ athlete disciplines**.

### Project Associate | Multi-Modal Perception & Satellite Detection

May 2023 – May 2025

Indian Institute of Science (IISc) & ISRO Space Applications Centre

Bengaluru, India

- Constructed an automated Python pipeline processing **100+ GB/day** of multispectral HDF5 imagery (~14,000 frames); two-stage PyTorch detector assessed across **5+ sensor bands** (mAP, RMSE), lifting throughput by **30%**.
- Slashed inference latency **from ~3–4 hrs to <1 sec** via pruning, INT8 quantization, and parallelized preprocessing; co-authored **IEEE IGARSS 2025** paper on satellite cyclone detection.

### Research Intern | Generative Models & Anomaly Detection

Jan. 2023 – May 2023

Georgia Institute of Technology (Remote)

Atlanta, GA

- Synthesized threat datasets with **GANs** (+15% robustness); trained autoencoder-based detectors (+12% recall); combined with a supervised ensemble to reach **95% accuracy**; streamlined preprocessing, cutting time by **25%**.

### Data Science Intern | ML Pipelines & Analytics

Jun. 2022 – Mar. 2023

Deloitte Haskins and Sells LLP

Bengaluru, India

- Optimized Spark ETL pipelines (saved **15 hrs/week**); overhauled audit testing (**40%** time reduction); deployed **85%** of Power BI dashboards and exception-analysis tools adopted across **4+ audit teams**.

## PROJECTS

### 3D Human Reconstruction | PyTorch, CUDA, NeRF, Gaussian Splatting, SMPL, COLMAP

Aug. 2025 – Present

- Researching **monocular 3D human body reconstruction** under Prof. Cheng Zhang; building a neural rendering pipeline to convert single-view RGB video into geometrically accurate, pose-driven 3D meshes.
- Evaluating reconstruction quality against SOTA baselines (PSNR, SSIM, LPIPS); modeling human-object interactions with SMPL body priors targeting sub-centimeter accuracy on ZJU-MoCap and Human3.6M.

### Pose4DGS | PyTorch, CUDA, 4D Gaussian Splatting, DINOv2, SMPL, W&B

Jan. 2026 – Present

- Implemented **Pose-Conditioned 4D Gaussian Splatting** for pose-driven avatars; added Pose-Space Augmentation (SMPL interpolation/extrapolation) and **DINOv2 (ViT-B/14)** feature supervision as zero-overhead training-only contributions.
- Formulated a **3-tier evaluation protocol** and **2×2 ablation grid**; logged **200+ runs** in W&B; co-authored a 6–8 page report with a 3-person cross-functional team.

### Recurrent Attention Model (RAM) | PyTorch, REINFORCE | GitHub

Aug. 2025 – Dec. 2025

- Reproduced **RAM** in PyTorch for MNIST classification via sequential glimpses (REINFORCE), reaching **96.64% accuracy** with **6 glimpses (~49% of pixels)**—matching full-image CNN baselines at a fraction of the compute.
- Deployed a learned value baseline to reduce gradient variance; ablated glimpse count and patch resolution; visualized attention trajectories confirming digit-edge fixation patterns.

### CycloAI | PyTorch, CUDA, GDAL, Rasterio, Spark, Dask, Docker, W&B

May 2023 – May 2025

- Co-built an end-to-end **multi-sensor perception pipeline** on **100+ GB/day** of satellite data; two-stage detector with self-supervised pre-training cut **inference from ~3–4 hrs to <1 sec**; raised minority-class recall **12%** via focal loss and SMOTE; published at **IEEE IGARSS 2025**.

## PUBLICATIONS

- Comparative Analysis of Deep Learning Models for Cyclone Detection and Localisation on Infrared Satellite Images.** Agrawal\*, Mohapatra\* et al. *IEEE IGARSS 2025*, Brisbane, Australia.
- Generative Adversarial Network Based Approach towards Synthetically Generating Insider Threat Scenarios.** Mohapatra, Phukan, Madiseti. *J. Software Engineering & Applications*, 16(11):586–604.