

TARUN SADARLA

Denton, TX 76201 • 940.765.8671 • tarunsadarla26@gmail.com

[LinkedIn](#) • [GitHub](#) • [Kaggle](#)

CAREER PROFILE

Medical Imaging AI specialist completing an MS in Artificial Intelligence (Biomedical Concentration, GPA: 4.00) at the University of North Texas, graduating May 2026. Focused on applying deep learning to clinical imaging challenges — including radiology (DICOM/MRI), ultrasound segmentation, histopathology, and neuroimaging. Experienced across the full ML lifecycle: data preprocessing, model development (CNNs, Vision Transformers, U-Nets), evaluation, and deployment. Actively seeking roles in healthcare AI, particularly in radiology and medical image analysis.

EDUCATION

University of North Texas | Denton, Texas

Expected: May 2026

Master of Science – Artificial Intelligence | GPA: 4.00 | Concentration: Biomedical

Relevant Coursework: AI for Wearables and Healthcare, Applications of AI in Health, Deep Learning, Machine Learning, Adv Pattern Recognition, Feature Engineering, Fundamentals of AI, Software Development for AI

Manipal Institute of Technology | Manipal, Karnataka, India

December 2023

Bachelor of Technology – Computer Science & Engineering | GPA: 3.29 | Minor: Data Analytics

Relevant Coursework: Deep Learning, Computer Vision, Big Data Analytics, Data Warehouse and Data Mining

TECHNICAL SKILLS

Medical Imaging & Vision: DICOM, MRI/Radiology image analysis, U-Net segmentation, OpenCV, PIL, CNNs, Vision Transformers

ML & Deep Learning: PyTorch, TensorFlow, Keras, Scikit-learn, Minitorch

Programming: Python, SQL, MATLAB, C

Cloud & Deployment: AWS, Docker, Flask, FastAPI

Big Data & Tools: Hadoop, Hive, PySpark, Cassandra, Git, Anaconda

ENGINEERING PROJECTS

Cine-Loop Based Fetal Head Measurement Framework (Group)

Aug 2025 – Dec 2025

- Designed a cine-loop segmentation framework extending static ultrasound analysis to temporal sequences using Residual U-Net architecture, targeting robustness to probe motion, speckle noise, and boundary instability.
- Replicated and validated a published Residual U-Net baseline for automated fetal head circumference estimation on the HC18 dataset, implementing the full preprocessing, segmentation, and evaluation pipeline.
- Built an end-to-end evaluation pipeline combining segmentation metrics (Dice, IoU) with biometric accuracy (HC MAE), enabling systematic assessment under realistic imaging conditions.
- Benchmarked against published results and a resource-constrained cine-loop model, analyzing the impact of data quality, resolution, and temporal variability on performance.

Deployable Histopathologic Cancer Detection System (Group)

Aug 2024 – Sep 2024

- Developed a custom 5-layer CNN with batch normalization, Leaky ReLU, max pooling, and dropout for binary classification of cancerous vs. non-cancerous pathology images.
- Achieved ROC-AUC of 0.81 and Precision-Recall AUC of 0.73; reduced false negatives through dataset scaling and systematic hyperparameter tuning.
- Served as one of two model developers in a team of five, leading deep learning design, implementation, and fine-tuning.

Detection of Autism Spectrum Disorder Using Deep Learning (Individual) *Jul 2023 – Nov 2023*

- Developed CNN and Vision Transformer (ViT) models for ASD detection from neuroimaging data, reducing subjectivity in clinical evaluations.
- Achieved 93–97% accuracy with CNN model; systematically diagnosed ViT limitations under resource constraints, informing architecture trade-offs for medical imaging tasks.
- Compared three optimizers across both architectures, producing a structured analysis of convergence behavior on clinical neuroimaging data.

WECARE: Wearable Emergency Cardiac and Fall Response System (Group) *November 2025*

- Designed an on-device real-time AI framework for closed-loop fall and cardiac event detection, fusing ECG and IMU signals without cloud dependency.
- Built parallel 1D CNN pipelines for fall detection (MobiFall) and arrhythmia detection (MIT-BIH ECG); achieved $F1 = 0.81$ and $F1 = 0.99$ respectively with edge inference latency under 40 ms.
- Implemented preprocessing pipelines including bandpass filtering, wavelet denoising, windowed segmentation, and class-balancing for severe imbalance scenarios.

Channel Redundancy Detection and Structured Pruning in CNNs (Individual) *Aug 2025 – Dec 2025*

- Designed a “Tier-2 Hybrid Crossover” pruning framework for VGG16, replacing binary filter elimination with regression-based weight synthesis to preserve feature information under aggressive compression.
- Achieved 2x channel compression (352 vs. 182 channels in deep layers) while improving Top-1 accuracy by +0.37% and accelerating convergence on CelebA dataset.
- Engineered dynamic safety guards and iterative rollback mechanisms enforcing global (2%) and per-class accuracy constraints throughout optimization.

Multimodal Sentiment and Emotion Classification (Video & Audio) (Group) *Mar 2025 – Apr 2025*

- Implemented unimodal and multimodal deep learning pipelines for emotion/sentiment recognition using the MELD dataset, combining CNN-based visual features with BiLSTM audio modeling.
- Conducted early-fusion multimodal experiments comparing baseline and pretrained architectures (VGG-16, deeper BiLSTM); analyzed class imbalance impact on minority emotion recognition.

Big Data-Driven Customer Segmentation in E-Commerce (Group) *Feb 2025 – Apr 2025*

- Built an end-to-end analytics pipeline on AWS (S3, EMR, Hadoop/Hive, PySpark, Spark MLlib) to preprocess and segment large-scale transactional data.
- Engineered RFM features and applied K-Means and Gaussian Mixture Models; evaluated clustering via silhouette metrics and PCA-based visualization for marketing insights.