

HARSHITH CHENNUPATI

Baltimore, MD | harshith@chennupatis.com | +1 443 615 9830 | [LinkedIn](#) | [GitHub](#)

SUMMARY

ML and systems engineer who builds real-time multimodal AI pipelines and ships them into deployed environments. Designed and adapted ControlNet-conditioned 3D latent diffusion models, transformer-based sequence classifiers over large physiological time-series, and GPU-accelerated 3D computer-vision pipelines integrating robotic and multi-sensor hardware. Strong foundations in PyTorch, CUDA, ONNX/TensorRT, and modern generative AI, with an AI-native development workflow and a track record of decomposing fuzzy problems into shippable components.

EDUCATION

Johns Hopkins University, MSE Electrical & Computer Engineering *Baltimore, MD | 05/2026*

GPA: 3.8/4.0 | Concentration: AI, Generative Deep Learning, 3D Computer Vision, and Medical Device Design

SRM Institute of Science & Technology, B.Tech Electronics & Communication *Chennai, India | 06/2024*

GPA: 3.76/4.0

Relevant Coursework: Deep Learning, ML for Medical Applications, Computer Integrated Surgery, ML for Signal Processing, Convex Optimization, Reinforcement Learning, DSP.

TECHNICAL SKILLS

- **Languages:** Python, C/C++, CUDA, Bash, MATLAB, SQL, JavaScript, R, LaTeX.
- **ML & Generative AI:** PyTorch, TensorFlow, Hugging Face, Diffusers, MONAI, scikit-learn, NumPy, SciPy, pandas, statsmodels; latent diffusion, ControlNet, VAEs, GANs, transformers, LLM fine-tuning and prompting, knowledge distillation, self-supervised pretraining, statistical modeling.
- **3D & Geometric Computer Vision:** 3D point cloud processing (ICP, feature-based), 3D-2D registration, DRR generation, FDK reconstruction, volumetric CNNs, mesh and volume resampling; ITK, VTK, 3D Slicer with Qt module development, Open3D, OpenCV.
- **Systems & Performance:** CUDA-based 3D imaging library, ONNX, TensorRT, multi-GPU training (RTX 3090, 2080, A100), distributed data loading, performance profiling, real-time inference pipelines.
- **MLOps & Cloud:** AWS (EC2, Lambda, S3), RunPod and remote multi-GPU training, SLURM-style HPC, Docker, Git, Weights & Biases, experiment tracking, dataset versioning, evaluation harnesses (LPIPS, SSIM, PSNR, FID, AUROC, calibration), pytest, CI.

RELEVANT EXPERIENCE

Research Assistant, ISTAR Lab, Johns Hopkins University *Baltimore, MD | May 2025 to Present*

- Built MAISI-US, a ControlNet-conditioned 3D latent diffusion pipeline adapting NVIDIA's MAISI foundation model to cross-modality ultrasound-to-CT synthesis. First reported adaptation of MAISI to a non-CT input modality; trained only the ControlNet branch on 9,068 paired US-CT volumes for low-cost domain transfer while keeping the VAE and diffusion backbone frozen.
- Designed a multi-metric evaluation protocol (PSNR, mutual information, gradient orientation agreement, LPIPS) against paired ground truth, with parameter sweeps on inference steps, ControlNet scale, and noise factor for reproducible characterization of model behavior.
- Built GPU-accelerated 3D vision pipelines (volumetric resampling, 3D-2D registration, DRR generation) on a CUDA-based imaging library; implemented pose regression with a custom Res-UNet for real-time 2D-in-3D visualization.
- Led multimodal data acquisition integrating Clarius and Alpinion ultrasound probes, Medtronic O-arm CBCT, NDI optical tracking, and scripted Universal Robot (UR3) sweeps to produce a paired calibrated 3D dataset for training and quantitative evaluation.
- Shipped Qt-based tooling for 3D Slicer to orchestrate deep-learning inference, putting research models directly into the hands of clinicians and researchers.

Algorithms Engineer, Neurava *Baltimore, MD | Jan 2026 to Present*

- Built a PSG-teacher to IMU-student knowledge distillation pipeline for wearable sleep staging on the DREAMT corpus; benchmarked classical biosignal baselines (Cole-Kripke, vanHees) against transformer-based sequence models with a rigorous evaluation harness across paired biological corpora.

- Developed IMU-based human pose estimation with clustering (K-means, GMM, HDBSCAN) on bicep-mounted sensor streams; authored the ground-truth labeling protocol for posture-class validation.
- Established MLOps practices for the team: experiment tracking, dataset versioning, reproducible training configs, and evaluation dashboards for fast iteration.

Biomedical Signal Processing Intern, Vigo

Remote | Spring 2025

- Designed ECG-based respiratory rate estimation using CEEMDAN, discrete wavelet transforms, and peak detection, achieving 2.5 bpm MAE on the BIDMC dataset.
- Implemented and unit-tested accelerometer-based posture classifiers for a wearable monitoring product.

SELECTED PROJECTS

3D Point-Cloud Registration & Surgical Navigation Algorithms (C++ / Python)

Fall 2024

- Implemented 3D point-cloud registration, calibration, error-modeling, and distortion-correction modules; cut ICP runtime 40% with Bounding Volume Hierarchies and reduced EM-tracking positional error 30%. Authored per-module unit tests and verification documentation.

Pneumothorax Segmentation from Chest X-rays (PyTorch)

Spring 2025

- Trained a multi-task U-Net with residual blocks on the SIIM-ACR pneumothorax dataset; achieved IoU 0.65 and 80% accuracy with a custom class-imbalance loss. Evaluated on Dice and IoU and analyzed failure modes across subgroups to quantify model confidence.

Deep Unrolled Networks for MRI Reconstruction (PyTorch)

Spring 2025

- Implemented ISTA-Net and ADMM-Net for accelerated MRI reconstruction from 25% undersampled k-space; extended ADMM-Net with Squeeze-and-Excitation blocks for 1.14x speedup at 24 dB PSNR.

Noise-Resilient Multimodal Emotion Recognition (NLP + Speech)

Fall 2024

- Improved noisy-audio classification accuracy from 18% to 51% using MetricGAN+ for denoising and Wav2Vec2 for robust speech features; gained an additional 20% by adding HuBERT and BERT multimodal fusion on the MELD dataset.

ADDITIONAL

First-author manuscript in preparation: “Cross-Modality Generalizability of Medical Image Synthesis Models”, On ultrasound-to-CT synthesis via ControlNet adaptation of the MAISI foundation model (ISTAR Lab, Johns Hopkins University).

Former national-level basketball player. Passionate about fitness and exercise science, with a strong interest in wearable technology and its impact on health monitoring and performance optimization.