HARSHITH CHENNUPATI

Baltimore, MD | harshith@chennupatis.com | LinkedIn | GitHub

EDUCATION

Johns Hopkins University, MSE Electrical & Computer Engineering — Baltimore, MD | May 2026
GPA: 3.8/4.0. Coursework: Deep Learning, ML for Medical Applications, Medical Imaging Systems, ML for Signal Processing, Computer Integrated Surgery, Convex Optimization, DSP.

SRM Institute of Science & Technology, B.Tech Electronics & Communication — Chennai, India | June 2024
GPA: 3.76/4.0

EXPERIENCE

Algorithms & Software Engineer, Neurava (full-time) — Baltimore, MD | Spring 2026 – Present
Wearable seizure-monitoring system for epilepsy (arm and neck devices) on an FDA regulatory track.

ML and clinical algorithms

- Developed the deep-learning sleep/wake algorithm for a clinical epilepsy wearable, reaching over 80% accuracy on held-out subjects; accepted as an AES 2026 poster.
- Took the sleep algorithm from research to production, where it supports analysts' day-to-day clinical review and lays the groundwork for future sleep-algorithm work.

Device authentication, mobile and FDA documentation

- Own device authentication for an FDA-track medical wearable end to end, from firmware to the mobile apps, and shipped it across the company's apps.
- Author software and device-security documentation and verification evidence supporting the FDA submission.
- Own the iOS platform and core patient-facing features of Neurava's new mobile app, working across the firmware, cloud and regulatory teams.

Cloud platform and manufacturing

- Key engineer on Neurava's migration to a regulated medical-device cloud platform, implementing and shipping cloud services now in production.
- Building manufacturing and release software that lets the contract manufacturer produce devices with full traceability, and automating sensor verification for robotic production QA.

Research Assistant, I-STAR Lab, Johns Hopkins University — Baltimore, MD | May 2025 – Present

- Enabled real-time 2D-in-3D visualization for spine surgery by adapting NVIDIA's MAISI 3D diffusion model to ultrasound-to-CT (retrained ControlNet, froze VAE and diffusion backbone) and a Res-UNet pose-regression registrar.
- Designed the evaluation protocol (LPIPS, SSIM, PSNR) against paired cadaver ground truth, with reproducible multi-GPU training and Weights & Biases tracking.
- Led multimodal cadaver collection (ultrasound, cone-beam CT, optical tracking, robot arm) and shipped 3D Slicer modules for segmentation and synthesis.
- Proposed a fix to open-source PlusToolkit/PlusLib (PR #1238, 2025); maintainers shipped a similar fix upstream. Extended NVIDIA MAISI tutorial pipelines in a fork for ultrasound-to-CT experiments.

Teaching Assistant, Deep Learning for Medical Imaging, Johns Hopkins University — Baltimore, MD | Fall 2025

- Designed and graded assignments on CNNs, segmentation, and multimodal registration; mentored students on PyTorch training diagnostics.

Biomedical Signal Processing Intern, Vigo — Remote | Spring 2025

- Estimated respiratory rate from ECG (CEEMDAN, wavelets, peak detection; 2.5 bpm MAE) and shipped accelerometer posture classification with unit tests.

Intern, Life Signals — Bangalore, India | Jul – Aug 2022

- Supported ECG collection and testing for FDA-approved wireless biosensors in hospital settings.

PROJECTS

MedDataMCP (Python, MCP, uv) 2026

- Open-source MCP server that lets any MCP client run one federated, ranked search for medical and biosignal datasets across Zenodo, PhysioNet and NSRR (plus Hugging Face, OpenNeuro and Kaggle), including paired-modality queries such as PSG + accelerometry.
- Connector-based design with a shared normalized dataset schema, boundary-aware modality inference, TTL caching and per-source error isolation so one failing API never breaks a search; pytest suite.

ADOPA: COPD Lung-Sound Classifier (Python, librosa, scikit-learn)

- Exploratory research prototype that classifies COPD from respiratory sounds using librosa audio features and a Random Forest on the public ICBHI dataset; started as a class project with friends, grew into an investor pitch, and shipped with a browser front end and inference API.

Omarchy System Monitor (QML / Quickshell, Bash) 2026

- Low-overhead system dashboard for the Omarchy Linux bar (19 stars, 9 forks, ~3K installs): samples /proc and /sys directly with adaptive 2 s / 5 s polling, with no background daemon, helper processes or telemetry; GPUs it cannot read without spawning helpers are left out rather than reported as idle.

Surgical Navigation Algorithms (C++ / Python) Fall 2024

- Built 3D point-cloud registration, calibration, and distortion-correction for surgical navigation; cut ICP runtime 40% with bounding volume hierarchies and reduced EM-tracking positional error 30%.

Deep Unrolled Networks for MRI Reconstruction (PyTorch) Spring 2025

- Implemented ISTA-Net and ADMM-Net for sparse MRI reconstruction from 25% undersampled single-coil knee data; added Squeeze-and-Excitation blocks for 1.14× speedup at 24 dB PSNR.

Earlier projects

- ECG arrhythmia detection with CNN and LSTM models; noise-robust speech emotion recognition (MetricGAN+, Wav2Vec2, HuBERT + BERT audio-text fusion); software-defined-radio marine VHF scanner (GNU Radio, RTL-SDR, AIS).

PUBLICATIONS & PRESENTATIONS

- **Sleep/Wake Detection Using Deep Learning With a Bicep-Mounted Wearable Device.** H. Chennupati, T. Meyer, V. Ganesh, J. Shah. AES 2026 Annual Meeting, poster 1.102 (accepted).
- **Ultrasound-to-CT synthesis with MAISI for spine surgery** (I-STAR Lab, Johns Hopkins): manuscript in preparation.
- **SDR-based marine VHF scanner:** conference presentation, ICRETM.

SKILLS

Languages: Python, C/C++, CUDA, Swift, Kotlin, JavaScript/TypeScript, SQL, MATLAB, Bash

ML / DL: PyTorch, TensorFlow, MONAI, Hugging Face, Diffusers, scikit-learn, ONNX, TensorRT; CNNs, transformers, 3D latent diffusion (MAISI, ControlNet), U-Net / Res-UNet

Biosignals and imaging: IMU, EEG, ECG, EMG, PSG; CEEMDAN, wavelets; CT / MRI / ultrasound, multimodal registration, 3D Slicer, ITK, VTK; LPIPS / SSIM / PSNR evaluation

Device and cloud software: React Native, Swift/SwiftUI, Android, BLE, embedded firmware, AWS, SQL, GitHub Actions, Docker, uv, pytest, Git, W&B

Regulated development: FDA design-control documentation, verification evidence, requirements traceability, test protocols, IEC 62304 / V&V for medical-device software