Harshitha Govindaraju

PhD Candidate, Electrical & Computer Engineering — Rutgers, The State University of New Jersey

hg293@rutgers.edu · +1 (848) 252-1493 · New Brunswick, NJ · F-1 OPT/CPT — no sponsorship needed · Google Scholar · harshi.phd

PhD candidate (ECE, Rutgers) building deep-learning systems and hardware for biomedical imaging — image recovery, cell segmentation, and biomarker quantification on noisy, low-SNR biological data, shipped as open web tools other researchers can use. Work spans the full chain: low-noise analog front ends and COMSOL multiphysics device models through CNN, GAN, and Deep-Image-Prior pipelines, with GPU-accelerated CUDA on NERSC HPC clusters and ML inference on Xilinx FPGA-PYNQ and NVIDIA Jetson edge platforms. Two years of regulated medical-device R&D at Siemens Healthineers (SPECT). Eight peer-reviewed publications, including a lead classical-vs-DL benchmark in *The Analyst* (RSC).

EDUCATION

PhD, Electrical & Computer Engineering — Rutgers University 2020 – 2026 (expected)

Advisor: Prof. Umer Hassan · GPA 3.7/4.0 · Deep-learning–driven biosensing & biomarker quantification; hardware–algorithm co-design.

MS, Biomedical Engineering — Rutgers University 2018 – 2020

Computational imaging, deep learning, and quantitative measurement on noisy biological data.

B.Eng, Electronics & Communication Engineering — R.N.S. Institute of Technology, VTU (India) 2013 – 2017

GPA 4.0/4.0 · Best Outgoing Student.

RESEARCH & INDUSTRY EXPERIENCE

Graduate Research Assistant — Hassan Lab, Rutgers ECE 2020 – Present

Deep-learning systems and signal-recovery algorithms for translational biosensing.

- Authored **HIST-DIP**, a hybrid histogram-thresholding + Deep Image Prior framework recovering low-SNR fluorescence images with no paired training data (PSNR 15.6→27.1 dB); published in *The Analyst* (RSC) and released as a public web demo.
- Designed, trained, and benchmarked CNN, GAN, autoencoder, and DIP pipelines (PyTorch/TensorFlow) for biomarker quantification on fluorescence-microscopy and biosensor data; benchmarked against flow-cytometry ground truth on patient samples.
- Modeled sensor electromagnetics and noise in **COMSOL Multiphysics**; co-optimized device geometry, bias, and signal-chain parameters against system-level SNR/LoD targets, then validated against fabricated hardware.
- Designed and validated low-noise analog front-end signal-conditioning chains — bias, amplification, filtering, ADC interface — quantifying noise spectral density, dynamic range, drift, and repeatability.
- Released validated models as open, no-install web tools (denoisedip.vercel.app, cellquantai.com, graphene-hall-sensor.vercel.app, uflow.studio) so collaborators and clinicians can run inference and simulation directly.

R&D Assistant, SPECT R&D — Siemens Healthineers, USA 2019 – 2021

Deep-learning systems, GPU-scale image pipelines, and large-scale clinical data infrastructure in a regulated environment.

- Implemented **GAN-based unsupervised lesion detection** on SPECT nuclear images; benchmarked against classical iterative reconstruction with quantified gains in artifact suppression and lesion contrast.
- Built **GPU-accelerated CUDA pipelines** (C/C++, Python) for high-throughput signal and image processing on NVIDIA platforms and NERSC HPC clusters; profiled, optimized, validated against reference implementations.
- Engineered a **DICOM data lake** on MongoDB with MD5 integrity checks and a Python metadata-extraction plugin — turning raw detector scans into a queryable, reproducible cohort dataset (published).
- Developed multi-modal SPECT/CT fusion, reconstruction, denoising, registration, and segmentation with rigorous evaluation (SNR, accuracy, repeatability); deployed ML inference on Xilinx FPGA-PYNQ.

TECHNICAL SKILLS

Deep learning: PyTorch, TensorFlow, Keras; CNNs, GANs, autoencoders, Deep Image Priors, U-Net; transfer learning, self-supervised pretraining; GNNs (PyTorch Geometric).

Programming: Python (NumPy, SciPy, Pandas, scikit-learn, scikit-image, OpenCV), C/C++, MATLAB/Simulink, SQL, R, shell, Git.

Scaling & compute: CUDA/GPU pipelines, NERSC HPC clusters, mixed-precision, Docker, Linux; Xilinx FPGA-PYNQ & NVIDIA Jetson edge inference.

Imaging & signals: Denoising, segmentation, registration, multi-modal fusion (SPECT/CT); DICOM/NIfTI; FIR/IIR filtering, FFT, lock-in detection, calibration.

Modeling & hardware: COMSOL Multiphysics, low-noise analog front-end design, noise budgeting; oscilloscopes, lock-in amplifiers, SMUs, spectrum analyzers, LabVIEW.

SELECTED PUBLICATIONS — full list on Google Scholar

1. **HIST-DIP: Histogram Thresholding and Deep Image Priors–Assisted Smartphone-Based Fluorescence Microscopy Imaging.** *The Analyst, Royal Society of Chemistry (2025).*
2. **Machine-Learning–Enabled Leukocyte Quantification Using a Smartphone-Coupled 3D-Printed Microfluidic Biosensor.** *IEEE Access (2022).*
3. **Modeling Enhanced Detection Dynamics of Magnetic Particles with Halbach Array and Hall-Effect Sensing for Biomedical Applications.**
4. **Particle Quantification from a Smartphone-Based Biosensor Using Deep CNNs for Clinical Diagnosis.**
5. **A Modular Microscopic Smartphone Attachment for Imaging and Quantification of Multiple Fluorescent Probes Using Machine Learning.**
6. **DICOM Data Storage and Retrieval with MongoDB.**
7. *Under review, IEEE Sensors Journal:* Graphene Hall-Effect Biosensor for Multiplexed Biomarker Detection.

LIVE TOOLS & INTERACTIVE SIMULATORS

HIST-DIP — denoisedip.vercel.app · fluorescence recovery demo.

CellQuant AI — cellquantai.com · CNN cell quantification.

GHSL — graphene-hall-sensor.vercel.app · Hall-sensor simulator.

uFlow — uflow.studio · Three.js microfluidic Hall simulator.

TEACHING (RUTGERS ECE) & AWARDS

TA / lab instructor: DSP & Linear Systems (2021–24), Digital Electronics (2024–25), Principles of EE II (2021–22), Principles of CS (2019–20).

Rutgers Research Scholarship (2020–2025) · Best Outgoing Student, R.N.S. Institute of Technology (2017) · National Distinction in Drawing.