

SUDARSHAN ANAND

Atlanta, GA (willing to relocate)|asudarshan14@gmail.com|+1 (404) 453-5664|Linked In: www.linkedin.com/in/sanand315/

Portfolio: sudarshanand14.vercel.app | F-1 OPT (2026–2029)

EDUCATION

Georgia Institute of Technology

Atlanta, GA, USA

Master of Science in Computational Science and Engineering (CGPA: 4.0/4.0)

May 2026

Relevant Coursework: Data Science for Epidemiology, Machine Learning for Graphs, Industrial Math, Numerical Linear Algebra

Birla Institute of Technology and Science (BITS Pilani)

Pilani, Rajasthan, India

Master of Science in Mathematics and Bachelor of Engineering in Computer Science (CGPA: 9.4/10)

Aug 2024

TECHNICAL SKILLS

Systems & Infrastructure: HPC workflows (SLURM-based scheduling), distributed training pipelines, GCP (Cloud Run), FastAPI, PostgreSQL, Pinecone

Programming & Data Tools: Python, modern environment management (uv), SQL, Pandas, NumPy, Git/version control, benchmarking

Machine Learning / AI: PyTorch (distributed training, mixed precision), LLMs, Generative AI, inference optimization, RAG architectures, modular ML pipelines

PROFESSIONAL EXPERIENCE

AI Product Development Intern

Jun 2025 – Aug 2025

Rezolve.ai

Dublin, CA, USA [Remote]

- Shipped an MVP enterprise RAG system utilizing FastAPI, Pinecone, and PostgreSQL, achieving sub-20s retrieval latency for enterprise policy document queries, replacing previously manual workflows optimized embedding strategies and context-window design.
- Engineered a prefix-tree (trie)-based autocomplete system, enabling low-latency prefix search for both personalized and AI-curated suggestions and deployed automated global data pipeline updates using GCP Cloud Run to ensure high-availability context ingestion.
- Prototyped an LLM-based agent for automated alert triaging, generating structured incident-response workflows from noisy system alerts under varying prompt and context constraints.

AI Scientist Intern

Jan 2024 – Jun 2024

Qure.ai Technologies Pvt. Ltd.

Bangalore, Karnataka, India

- Optimized computer vision models utilizing PyTorch for clinical lung nodule detection, improving Pearson correlation metrics by ~45% relative to baseline models.
- Benchmarked temporal alignment consistency across longitudinal imaging datasets to validate model stability for production deployment.

RESEARCH AND PROGRAMMING PROJECTS

Multimodal Progression Tracking of Neurodegenerative Disorders, Georgia Institute of Technology

Aug 2025 – present

- Engineered a multimodal PyTorch pipeline for Parkinson's disease progression modeling, utilizing distributed SLURM clusters to orchestrate array jobs and process a cohort of over 3,000 high-dimensional 3D volumetric images (MRI scans).
- Fine-tuned a self-supervised vision encoder (SimCLR) with structured demographic feature fusion, optimizing GPU-accelerated training loops to handle complex, overlapping clinical phenotypes efficiently (macro F1 \approx 0.60).

Investigating importance of Patient Metadata in Public Health, Edith Cowan University

May 2025 – present

- Analyzed 1.2M unstructured patient records to quantify the impact of extreme metadata sparsity (94% missing rate) on downstream predictive modeling.
- Designed and executed probing experiments using state-of-the-art LLMs (Gemma-2, MedGemma) to rigorously evaluate model robustness, reasoning capabilities, and alignment under degraded context and missing-information regimes.

Samay: Modular ML Inference Framework, Georgia Institute of Technology

Jan 2025 – May 2025

- Developed [Samay](#), an open-source framework unifying 10 time-series foundation models, engineering standardized inference, fine-tuning, and evaluation workflows across diverse data modalities (34 stars, 7 forks).
- Built modular PyTorch 2.0+ pipelines to enable zero-shot evaluation and seamless task-specific adaptation (SFT), optimizing for fast execution across diverse data modalities.
- Architected and executed large-scale benchmarking protocols, validating model performance, scalability, and baseline metrics against unified evaluation standards.

HONORS & AWARDS

IEEE BHI 2025 Data Challenge Competition Winner

Oct 2025

- Engineered and developed an end-to-end AI/ML-based depression risk prediction system, contributing to a winning solution in the IEEE-sponsored competition (supported by NSF and Google)

IEEE BHI 2025 Young Professional NextGen Scholar (NSF, EMBS, Google Sponsored)

Sept 2025

SELECTED PUBLICATIONS

Foundation Models / ML Systems

- H. Kamarthi, S. Li, S. Anand, S. Das, A. Ramesvaran, A. Bhalerao and B.A. Prakash, "Samay: A Unified Pipeline Framework for Time-Series Foundation Models" VLDB 2026 (Demo track, submitted)