

Sainava Modak

+91-9678817677 | sainava.modak@gmail.com | linkedin.com/in/sainava-modak | github.com/Sainava

EDUCATION

B.Tech in Computer Science and Engineering

CGPA: 9.51/10

Amity University Uttar Pradesh

2023 – 2027

- Awarded a 100% merit scholarship for the full duration of the degree; ranked among the top 5% of the cohort.

Senior School Certificate Examination (Class XII, CBSE)

97%

Army Public School, Shillong

2020-2022

- Achieved the highest percentage among CBSE-affiliated schools in Meghalaya in the Science stream.

Secondary School Leaving Certificate (Class X, Meghalaya)

94.16%

St. Thomas Higher Secondary School

2008-2020

- Secured 2nd state rank among 50,081 candidates in the SSLC Examination.

TECHNICAL SKILLS

Programming: Python, Java, C++, TypeScript

Computer Science Fundamentals: Data Structures & Algorithms, Object-Oriented Programming, Operating Systems, Computer Networks, Database Management Systems

AI & Machine Learning: PyTorch, Graph Neural Networks (GNNs), LSTMs, Retrieval-Augmented Generation (RAG), OpenCV, YOLO

Backend & Infrastructure: Spring Boot, MySQL, PostgreSQL, Redis, Neo4j, Kafka, Docker, CI/CD

Cloud (AWS): Lambda, API Gateway, DynamoDB, AWS SAM

PROJECTS

AirForce: 3D Indoor Air-Quality Simulator | TypeScript, Three.js, AWS (Lambda, DynamoDB, SAM)

GitHub

- Built a simulator that shows where PM2.5 builds up and air goes stale inside a user-designed home.
- Modelled airflow with a multi-zone pressure network and a 3D voxel fluid solver, with a serverless AWS backend serving live PM2.5, wind and temperature.
- Translated simulated PM2.5 into health risk using India's AQI bands and WHO guidelines, and ranked household fixes with side-by-side before/after comparison.

LingoStream | Java, Spring Boot, Redis, PostgreSQL, WebSockets

GitHub

- Prevented conflicting edits during concurrent subtitle collaboration using Redisson distributed locks and JPA optimistic locking, while using Redis Cache-Aside to reduce repeated database reads.
- Implemented real-time communication using STOMP WebSockets and Redis Pub/Sub to propagate messages across backend instances, with stateless JWT authentication securing API mutations.
- Delivered a Docker-deployed backend supporting real-time collaborative subtitle editing across multiple instances.

CV-Graph-RAG: Structural Literature Synthesis | Python, Neo4j, Llama 3

GitHub

- Combined semantic search with structured relationships between research papers, methods, tasks, and datasets to improve literature synthesis beyond similarity-based retrieval alone.
- Developed a pipeline using Llama 3 for entity extraction, MiniLM embeddings for document retrieval, cross-encoder reranking, and Neo4j to organize relationships between research concepts.
- Generated grounded answers that connect relevant research papers with their methods, tasks, and datasets.

Tactical GNN Spatial Engine | Python, PyTorch, Graph Neural Networks

GitHub

- Used player-tracking data to study short-term possession retention around the ball carrier in football.
- Applied a GATv2 graph neural network to a player-interaction graph incorporating movement, distance, closing speed, and Time-To-Intercept (TTI).
- Predicted short-term possession retention and visualized tactical scenarios through an interactive dashboard.

EXPERIENCE

GeoGlyph Inc.

Remote, Japan

AI Engineer Intern

July 2026 – Present

- Developing real-time edge AI software to process video streams from multiple cameras on resource-constrained devices and convert detections into spatial coordinates using homography-based transformations.
- Building a privacy-preserving pipeline using PyTorch, YOLO, and OpenCV, transmitting only spatial coordinates to a central server for multi-camera integration and downstream analysis without sending raw video.

Open Health Systems Laboratory (OHS�)

Remote, India

Research Intern

May 2026 – June 2026

- Researched architectural approaches for a universal edge-driver system to enable communication between heterogeneous medical devices and the OHS� TELESYNERGY platform.
- Evaluated frameworks for FHIR-based data normalization and concurrent WebSocket transmission to support synchronized health-care data delivery as part of a research team.

OPEN SOURCE CONTRIBUTIONS

Hugging Face Transformers | Python | 2 Merged PRs

[View PRs](#)

- Added native masked Mean Squared Error (MSE) loss for Sapiens2 pose estimation, enabling training with masked targets.
- Added pose-estimation keypoint preprocessing to Sapiens2ImageProcessor, converting keypoints into UDP Gaussian heatmaps for model input.

CERN ROOT (TMVA-SOFIE) | C++ | 3 Merged PRs

[View PRs](#)

- Implemented ConvTranspose padding generation for explicitly specified output shapes, including asymmetric spatial matrices and underflow safeguards.
- Added constant-folding for Gemm when inputs are compile-time constants, eliminating redundant runtime inference computation.
- Added ONNX Asinh, Acosh, and Atanh operator support to TMVA-SOFIE, extending operator registration and generated inference code.

NVIDIA NeMo (Speech) | Python | 1 Merged PR

[View PR](#)

- Optimized the ASR BLEU metric by batching tokenizer updates, reducing CPU–GPU synchronization operations from $O(N)$ to $O(1)$ per batch.

PUBLICATIONS & RESEARCH

Deepfake Detection Using Optical Flow and Deep Residual Learning Models

[Read Paper](#)

International Journal of Computers and Applications, 2026

- Combined spatial features from RGB frames with temporal motion cues from TV-L1 optical flow, using a dual-stream MixedResNet feature extractor and a Bi-LSTM over 1024-dimensional features.
- Achieved 94.79% accuracy on FaceForensics++ and 99.34% on Celeb-DF (V2).

Compatibility-Guided Integration of Diffusion Acceleration Methods

Accepted at IEEE International Conference on Intelligent Sensing and Autonomous Systems for Sustainable Future (ISASF 2026)

- Developed an Interoperability Assessment Protocol and a hierarchical Failure Taxonomy to evaluate whether independently developed diffusion acceleration methods, such as token reduction and caching, can be combined.
- Implemented two hybrid FLUX.1-dev pipelines, achieving end-to-end speedups of $3.54\times$ and $1.12\times$ over their respective baselines.

HONOURS & AWARDS

Talent for Japan 2025 Winner

Video Submission

- Architected an AI Factory Compliance Assistant using YOLOv8, DeepSORT and MediaPipe to automate safety audits from 360-degree video, with a two-tier rule engine for universal and sector-specific checks.
- Awarded a fully funded week-long visit to Japan to present the solution to Japanese corporate officials and media.

CERTIFICATIONS

- **AWS Certified Cloud Practitioner**
- **Oracle Cloud Infrastructure 2025 Certified Generative AI Professional**
- **Introduction to Networks** — Cisco Networking Academy