

MITRA BRINDA MUKHERJEE

AI Engineer

+91 8240571283 | mitrabrindasjd24@gmail.com
linkedin.com/in/mitrabrindaa | github.com/mitrabrindaa



EDUCATION

B.Tech, Computer Science (AI & ML) — **Sister Nivedita University, Kolkata, India**

2022 – 2026

CGPA: 8.37

EXPERIENCE

Junior Software Development Engineer (AI) — **Steora Systems**

Feb 2026 – Present

Kolkata, India

- Designed and deployed dialogue-state-managed (DSM) AI chatbots for business automation workflows; built voice agents on the Retell AI platform and video agents on the Anam AI platform for customer-facing deployments.
- Built backend services (FastAPI, PostgreSQL, REST APIs) to support these AI deployments, running on AWS and Azure cloud infrastructure.
- Integrated CRM systems, email automation, and third-party APIs into AI-powered platforms in collaboration with engineering and product teams.

Research Intern, ML for Cybersecurity — **Prof. Marcus Botacin — Texas A&M University**

May – Jul 2025

College Station, Texas

- Built AutoSteamYARA, an AI-driven system for automated YARA rule generation, implementing the Apriori algorithm to extract detection patterns from malware binaries and security telemetry.
- Implemented adaptive stream-learning models (e.g. adaptive Random Forest via the River library) for online, continuously-updating malware pattern detection, replacing static batch retraining.
- Designed automated validation and benchmarking pipelines for reproducible experimentation, reliability analysis, and failure-case analysis.

PROJECTS

Code-Switching Voice Agent | *LiveKit, Python, Next.js, Sarvam AI, Groq*

- Built a real-time voice agent that handles Hindi/Bengali ↔ English code-switching — mixing languages mid-sentence the way Indian users actually speak — instead of forcing a single-language mode; users upload a PDF and query it by voice, with answers grounded only in that document.
- Architected the pipeline on LiveKit Agents (Python) with Sarvam STT in code-mix mode for multilingual transcription, Groq-hosted Llama 3.1 8B for low-latency inference, and Sarvam Bulbul TTS for natural Hindi voice output.
- Built a themed Next.js frontend (LiveKit React components) for document upload and live voice sessions, with session transcript logging for debugging and evaluation.

Deepfake Detection System | *PyTorch, FastAPI, Docker, Hugging Face*

- Built an ensemble deepfake classifier (XceptionNet + EfficientNet) achieving 96.72% test accuracy, 91.96% F1-score, and 97.57% specificity on a balanced 8,538-image subset of FaceForensics++ (FF++).
- Applied two-phase transfer learning with selective fine-tuning, MixUp augmentation, label smoothing, and Test-Time Augmentation (TTA) to improve generalization and inference robustness.
- Deployed as a Dockerized FastAPI inference service hosted on Hugging Face for scalable predictions.

VibeCheck Finder | *FastAPI, PostgreSQL, Next.js, Sentence-BERT, Vector Search*

- Built an AI-powered movie discovery platform enabling natural-language “vibe” search in place of keyword or title search, using Sentence-BERT embeddings with a hybrid retrieval and similarity-ranking pipeline.
- Designed a full-stack architecture (Python backend, Next.js frontend, PostgreSQL storage) with a modular, production-style recommendation-ranking system.

TECHNICAL SKILLS

Languages: Python, C++, C, SQL

AI / ML: PyTorch, Scikit-Learn, Computer Vision, NLP, Semantic & Vector Search, Stream Learning (River), Model Evaluation

Conversational AI: Dialogue State Management, LiveKit Agents, Retell AI, Anam AI, Sarvam AI

Backend: FastAPI, Flask, REST APIs, PostgreSQL, Supabase

Cloud & Tools: AWS, Azure, Docker, Hugging Face, Git/GitHub