

Agrannya Singh

📞 +91 9508901080 • ✉️ singh.agrannya@gmail.com • 🔗 [linkedin.com/in/agrannya](https://www.linkedin.com/in/agrannya) • 📷 [Agrannya-Singh](#) • 📄 [agrannyasingh](#)

EDUCATION

B.Tech in Computer Science & Engineering

Vellore Institute of Technology, Vellore

Aug 2023 – Jun 2027

Vellore, Tamil Nadu

Coursework: Data Structures & Algorithms, Operating Systems, DBMS, Machine Learning, Computer Networks

PROFESSIONAL EXPERIENCE

Samsung Research Institute Bangalore (SRI-B) – PRISM R&D

Bangalore (Remote)

Research Intern – LLM-Generated Intent States for Sequential Music Retrieval

Feb 2026 – Present

- **Sequential Recommendation Engine:** Built and trained a **3.5M+ parameter deep sequential model (PyTorch GRU with Self-Attention)** on 50M+ music interactions, accurately predicting next-track listening intent from streaming history across 156k+ tracks.
- **Contrastive Representation Learning:** Formulated a joint multi-task objective pairing **contrastive learning** with 4-class user intent classification, outperforming the **0.1655 semantic similarity baseline** with GPU mixed precision (**CUDA AMP**).
- **High-Speed Vector Search (IVFPQ):** Indexed 156k+ track embeddings using **FAISS vector quantization (IVFPQ)**, achieving **90% RAM compression** (1024B → 32B) while enabling sub-millisecond nearest-neighbor retrieval at scale.
- **High-Performance Data Pipeline:** Engineered an out-of-core **DuckDB pipeline** ingesting 156k+ tracks in **118 seconds** under 8 GB RAM limits; built automated multi-epoch checkpointing engines for fault-tolerant model training.

Novustrana

Hyderabad (Hybrid)

Full Stack Intern

Mar 2026 – Apr 2026

- **AI-Powered Employability Scoring:** Architected an automated candidate evaluation service orchestrating **Google Gemini 2.5-Flash** with Firebase Cloud Functions, serving a live production-ready POC evaluated across 100+ active user accounts.
- **Backend & Data Reliability:** Developed Express.js API routes with JWT authentication; optimized database query latency by **25%** via Firestore composite indexing and configured automated Point-in-Time Recovery (PITR) for zero data loss.

SambalPay Fintech Solutions

Gurgaon (Remote)

Software Engineering Intern

Feb 2026 – Mar 2026

- **Core Banking Infrastructure:** Containerized and deployed enterprise core banking backend services (**Apache Fineract**) via **Docker Compose** on **GCP**, tuning JVM heap and resource limits to guarantee 99.9% system availability.

Mokshapay

Remote

Founding Intern – Full Stack Engineer

Oct 2025 – Jan 2026

- **Cloud Infrastructure as Code (IaC):** Automated compliant GCP infrastructure using **Terraform**, provisioning isolated VPC networks, Cloud NAT, and reserved static IPs required for secure, production-grade KYC API vendor synchronization.
- **Database Architecture & Security:** Built transactional Loan Management schemas on PostgreSQL and Firestore; authored granular database security rules, preventing client-side data drift and access vulnerabilities.

ACADEMIC & PERSONAL PROJECTS

ScreenScout – Multistage Semantic Recommendation & MCP Agent Server

[🔗 HF Space](#) | [🔗 Source](#)

Tech: FastAPI, BAAI BGE Bi/Cross-Encoders, Model Context Protocol (MCP), Pinecone Serverless, SQLite, Azure

- **Multistage Semantic Funnel & Reranking:** Engineered a 4-stage recommendation pipeline over 45k+ films combining **BAAI bi-encoders** ('bge-base') for candidate search, Bayesian rating shrinkage, and **BAAI cross-encoders** ('bge-reranker-v2-m3') for reranking; cut P95 latency by **90.6% to 1.58s** with **0.784 NDCG@10** ranking relevance and zero hallucination.
- **Model Context Protocol (MCP) Agent Server:** Built an enterprise-ready **Model Context Protocol (MCP)** server over Server-Sent Events (SSE), enabling AI agents (e.g., Anthropic Claude) to query live catalog recommendations with a **100% tool completion rate** across 50 stress-test scenarios while mitigating cloud reverse-proxy drops.

TuneTrace – Real-Time Semantic Music Discovery Microservice

[🔗 Live](#) | [🔗 Source](#)

Tech: Next.js, FastAPI, PostgreSQL (pgvector), SentenceTransformers, Redis, Docker, Supabase

- **Real-Time Dynamic Taste Profiling:** Engineered a sub-1.2ms taste profiler with **recency-weighted decay**, detecting and adapting to listener mood pivots in **<48 seconds** to eliminate collaborative filtering recommendation inertia.
- **Vector Similarity Search:** Implemented an end-to-end semantic retrieval pipeline using **SentenceTransformers** and **PostgreSQL (pgvector)**, retrieving matching tracks in **<45ms** via cosine distance queries.
- **Fault-Tolerant Microservice Architecture:** Designed a 4-tier fallback cascade (Semantic → Heuristic → YouTube API → Global Charts) guaranteeing **100% request fulfillment** (<85ms SLA) with a **24h Redis anti-repetition filter**.

Othello Dojo – Neural Network Strategy Agent

[🔗 Live](#) | [🔗 Source](#) | [🔗 Notebook](#)

Tech: Next.js, Node.js, Tailwind CSS, Python, MongoDB, Flask, TensorFlow

- **Game AI Engine:** Developed an intelligent strategy agent pairing deep convolutional neural networks (**ResNet architecture**) with optimized **Alpha-Beta Minimax search**, securing a 50% win rate against benchmark heuristic engines.

TECHNICAL SKILLS

- **AI, ML & Agentic Systems:** PyTorch, BAAI BGE (Bi/Cross-Encoders), Model Context Protocol (MCP / SSE), Vertex AI, Vector Search (FAISS, pgvector, Pinecone), Large Language Models (LLMs), SentenceTransformers, Contrastive Learning, TensorFlow
- **Databases & Storage:** PostgreSQL (pgvector), Pinecone Serverless (HNSW), DuckDB (OLAP), Redis, MongoDB Atlas, Cloud Firestore, SQLite
- **Frameworks & Cloud:** FastAPI, Next.js, React, Node.js, Express.js, Flask, Docker, GCP (Vertex AI, Cloud Run), AWS (EC2, S3), Terraform, GitHub Actions
- **Programming Languages:** Python, Java, TypeScript, JavaScript, SQL, HTML/CSS, R

ACHIEVEMENTS & CERTIFICATIONS

- **Academic Honors:** Secured **State Rank 22** in the National Talent Search Examination (NTSE Stage 1).
- **Professional Certifications:** [Claude with Vertex AI](#) (Google Cloud) · [Atlas Vector Search for RAG](#) (MongoDB) · [Generative AI](#) (IBM) · [DevOps and SRE](#) (Linux Foundation) · [Cloud Computing](#) (IBM)