

Abdul Ahad

Email: 099ahadabdul@gmail.com | Profile: <https://990aa.github.io/> | LinkedIn: <https://www.linkedin.com/in/01ahad/> |
Github: <https://github.com/990aa> | HuggingFace: <https://huggingface.co/a-01a>

Personal Profile

Dual-degree undergraduate grounded in advanced mathematics, with parallel focus on AI research and software engineering. My work spans deep learning, algorithmic optimization, and mathematical programming, alongside full-stack and cross-platform application development.

Education

Manipal Academy of Higher Education | B.Tech. Computer Science and Engineering (AI/ML Specialization) |

2024-2028 | CGPA: 9.14 of 10

Linear Algebra, Differential Equations, Numerical Methods, Discrete Mathematics, Object Oriented Software Development, Data Analytics, Database Systems, Automata Theory & Compiler Design, Quantum Computing, OS, Computer Networks

Indian Institute of Technology Madras | B.S. Data Science and Application | 2024-2028 | CGPA: 9.39 of 10

Probability & Optimization, Statistics, Data Structures & Algorithms, Machine learning, Deep Learning, Generative AI, Application Development, Tools in Data Science

Technical Skills

Languages: Python, Java, SQL, C, Typescript, Dart, Rust

AI & ML: PyTorch, TensorFlow, Scikit-learn, XGBoost, CatBoost, LightGBM, Transformers, OpenCV, YOLO, MediaPipe, Spacy

GenAI & LLM: LangChain, LlamaIndex, LangGraph, Unsloth, llama.cpp, PEFT/LoRA

Data Analysis: Polars, DuckDB, NumPy, SciPy, Matplotlib, Seaborn, Plotly, Pandas

Cloud & Backend: FastAPI, Django, Flask, Node.js, Uvicorn/Gunicorn, Celery, Redis, AWS

Frontend & Mobile Development: Next.js, Flutter, Tauri, JavaFX

DevOps & Databases: OracleDB, PostgreSQL, SQLite, MongoDB, Gradle, Docker, Git/GitHub/GitLab

Core Projects

[Kivixa Productivity Workspace](#) | Flutter, Dart, Rust, llama.cpp

- Architected a privacy-first, cross-platform productivity system with a Rust-native AI engine supporting multi-model SLM inference (Phi-4, Qwen 3.5, Gemma) with automatic task-based model routing, GPU acceleration via Vulkan/Metal, and a full Model Context Protocol (MCP) implementation for sandboxed AI-driven file operations — all fully offline.

- Built a Rust-backed audio intelligence pipeline integrating Whisper STT with real-time word-level timestamps, Kokoro neural TTS, voice activity detection, and semantic audio indexing; alongside a vector database with local embeddings for semantic search and automatic note clustering.

- Engineered version control using SHA-256 content-addressable blob storage with automatic snapshots, a Lua-scriptable plugin system with a full programmatic notes API, an interactive knowledge graph with force-directed simulation, and a split-screen editor supporting handwritten, markdown, and PDF formats.

[Unsupervised Cipher Cracking](#) | Python, NumPy, MCMC, HMM, Genetic Algorithms

- Implemented three independent solvers — Metropolis-Hastings MCMC with simulated annealing, Baum-Welch HMM/EM, and a tournament-selection genetic algorithm — over a shared quadgram language model trained on 6M+ characters, with vectorised NumPy decryption in the hot loop achieving ~100k scores/second.

- Built a statistical cipher-type detector using Index of Coincidence, Kasiski examination, entropy, and bigram repeat rate; and a phase-transition analyser that empirically maps minimum ciphertext length against decryption success rate, directly connecting to the theoretical unicity distance.
 - Designed a benchmarking and convergence diagnostics framework comparing all three solvers on identical test cases across five cipher families, with multi-chain consensus key reconstruction, per-letter confidence scoring, and publication-quality diagnostic plots.
- [NovelCrafter: Fine-Tuned Literary LLM](#) | PyTorch, Hugging Face, PEFT/LoRA
- Fine-tuned Meta's Llama 3.1 for long-form literary generation, using LoRA to adapt attention mechanisms while maintaining base model integrity.
 - Designed an incremental training pipeline for literary text to capture stylistic nuances without catastrophic forgetting, publishing the optimized adapters to the Hugging Face hub.
- [Hospital Operations System](#) | Flask, Vue.js, Celery, Redis, Docker
- Built a role-separated web platform across two integrated runtime modules — hospital operations and blood bank management — behind a shared Flask-Security authentication layer with four distinct RBAC roles, deployed via Docker Compose with Gunicorn, tini for correct PID handling, and a non-root production runtime.
 - Engineered a blood bank allocation engine using SQL triggers, compatibility-based matching logic, and predictive shortage alerting with a full forensic audit trail; donor workflows support whole blood and component-split donations with inventory state tracked across requests.
 - Implemented async task infrastructure with Celery workers and Beat scheduling for automated appointment reminders, monthly doctor activity reports, and CSV exports via PyMuPDF and Pandas — all brokered through Redis with a separate caching layer to minimise database hits.
- [Phantom Local AI Overlay Assistant](#) | Rust, Python, TypeScript, Dart, Kotlin
- Built a cross-platform AI assistant that idles at under 15 MB RAM using a multi-process architecture — a Rust watcher handling OS-level hotkeys, UIAutomation context extraction, and IPC via Named Pipes, with a Python inference engine loading GGUF models on-demand via llama-cpp and a Tauri/React floating overlay as the UI.
 - Implemented an Android counterpart in Flutter and Kotlin using an Accessibility Service for active-window context traversal and a native llama.cpp bridge for on-device inference, distributed as F-Droid-compatible APKs alongside a Windows WinGet package.
 - Designed a style distillation system that automatically extracts personalised writing rulebooks from outgoing message history, applied at inference time to mirror the user's tone without fine-tuning.

Certifications

[AWS Academy Graduate - Machine Learning Foundation](#): Amazon Sagemaker, Rekognition, Comprehend, Forecast, Lex, Kendra, Fraud Detection

[RAG & Agentic AI Professional](#): Certified in LangChain, LlamaIndex, and Agentic Frameworks (CrewAI, AutoGen).

[MLOps Specialization](#): Expertise in Azure ML, MLFlow, and Hugging Face deployment pipelines.