Abdul Ahad

Phone:+919141685085 | 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

Education

Manipal Academy of Higher Education | B.Tech. Computer Science and Engineering | *2024-2028* | *CGPA: 9.14/10*

Multivariable Calculus & Numerical Methods, Linear Algebra, Discrete Mathematics, Automata Theory & Compiler Design, DevOps & Cloud Computing

Indian Institute of Technology Madras | B.S. Data Science and Application | *2024-2028* | *CGPA: 9.5/10*

Deep Learning, Generative AI, Probability & Optimization, Algorithm Design, Machine learning, Tools in Data Science, Application Development

Technical Skills

Languages: Python, Rust, TypeScript, Dart, SQL
Backend & Cloud: FastAPI, Flask, Node.js, Celery, Redis, AWS, Docker
Frontend & Cross-Platform: Vue.js, Flutter, Tauri
Databases & Observability: PostgreSQL, SQLite, MongoDB, Prometheus, Grafana
Applied AI: PyTorch, Hugging Face, llama.cpp, Scikit-learn, LangChain, NumPy

Core Projects

Kivixa Productivity Workspace | Flutter, Dart, Rust, llama.cpp

• Architected an offline, cross-platform workspace (distributed via F-Droid and winget) driven by a Rust-native AI engine. Engineered multi-model LLM inference with Vulkan/Metal GPU acceleration and integrated a sandboxed Model Context Protocol (MCP) for secure, autonomous file operations.

• Built a Rust-backed audio intelligence pipeline featuring Whisper STT for real-time word-level transcription, Kokoro neural TTS, and voice activity detection; integrated alongside a local vector database for semantic search and a high-performance math engine for multivariable calculus and hypothesis testing.

• Implemented Git-like version control utilizing SHA-256 content-addressable blob storage with automatic snapshots. Designed a Lua 5.3 scriptable plugin architecture with a programmatic API, an interactive force-directed knowledge graph, and parallel state-driven orchestration for chained routines.

Hospital Operations System | Flask, Vue.js, Celery, Redis, Docker, Prometheus

• Architected a dual-module hospital and blood bank platform deployed via Docker Compose with a non-root Gunicorn/tini runtime. Secured the API with dual Session/JWT authentication, Argon2 hashing, and strict validation (Pydantic v2, Flask-Limiter, Talisman CSP).

• Engineered an SQL-heavy allocation engine with trigger-backed forensic auditing and compatibility matching. Implemented a Redis/Celery async pipeline for scheduled workflows and PDF/Pandas CSV exports, supported by SQLite WAL tuning to handle concurrent booking contention.

• Integrated production-grade observability featuring structured JSON logging (structlog) with cross-request X-Request-ID tracing, RFC 7807 error contracts, dependency health probes, and custom Prometheus business metrics visualized in Grafana.

Cryptex: Unsupervised Cryptanalysis Engine | Python, NumPy, Numba, MCMC, HMM

• Engineered a multi-language cryptanalysis suite supporting 7 classical cipher families, implementing adaptive MCMC (with parallel tempering), Baum-Welch HMM/EM, and KPA-constrained solvers optimized over Kneser-Ney smoothed 5-gram language models.

• Formulated a statistical cipher-type detector (leveraging Index of Coincidence, Kasiski examination, and entropy) alongside a phase-transition analyzer that empirically maps ciphertext length against decryption success rates to validate theoretical unicity distances.
• Executed a rigorous 130-run empirical benchmarking framework, achieving a 100% success rate and 0.0 Symbol Error Rate (SER) on substitution and transposition ciphers by accelerating the evaluation hot loop with Numba JIT-compiled tensor operations.

Universal Adversarial Cloak | PyTorch, Hugging Face, FaceNet, CLIP

• Engineered an adversarial optimization engine using Projected Gradient Descent (PGD) to mathematically cloak biometric and semantic identities against FaceNet and CLIP architectures.
• Enforced strict L-infinity norm constraints within the joint loss function to guarantee human-imperceptible perturbations, achieving a structural similarity index (SSIM) of > 0.98 across all cloaked outputs.
• Validated attack efficacy via an automated offline benchmarking pipeline, proving high black-box transferability by successfully blinding industry-standard oracle models (ArcFace, CLIP ViT-L/14) and optimizing Mean Residual Similarity (MRS) through comprehensive ablation studies.

GPU-Accelerated Geometric Image Reconstruction | PyTorch, NumPy, CUDA, OpenCV

• Architected a GPU-accelerated sequential hill-climbing engine to iteratively reconstruct high-fidelity photographic images using geometric primitives, employing a coarse-to-fine multi-stage schedule to capture both foundational structures and micro-details.
• Formulated an analytic closed-form color solver and integrated Sobel-based gradient maps for structure-aware shape routing, completely eliminating stochastic color search loops and maximizing algorithmic convergence speed.
• Executed rigorous empirical ablation studies across 11 architectural variants, achieving a peak Structural Similarity (SSIM) of 0.6059 and PSNR of 17.96 dB within a strict 1-minute compute budget, significantly outperforming standard evolutionary baselines.

Certifications

AWS Academy Graduate - Machine Learning Foundation**:** Expertise in integrating cloud-native AWS AI services (SageMaker, Comprehend, Fraud Detection) into application backends.
MLOps Specialization**:** Certified in Azure ML, MLFlow, and Hugging Face inference pipelines.
RAG & Agentic AI Professional**:** Certified in LangChain, LlamaIndex, and orchestrating multi-agent frameworks (CrewAI) for software applications.