Raghav Pathak

AI Engineering | LLM Applications | Retrieval-Augmented Systems | Applied Cryptography

Jaipur / Noida, India | +91 96504 88877 | raghavpathak0130@gmail.com

github.com/raghavpathak30 | linkedin.com/in/raghav-pathak-b14b46282

EDUCATION

The LNM Institute of Information Technology (LNMIIT), Jaipur — B.Tech, Computer Science and Engineering 2023 – 2027

EXPERIENCE

Omni Infoword Pvt. Ltd. — Cybersecurity Intern May – Jul 2026 (7 weeks)

- Ran an OWASP Top 10 test pass over a pre-release internal application, uncovering publicly reachable staging pages and exposed sensitive files before production release.
- Traced every finding to the developers who wrote the code, walked them through the exploit mechanism until each issue closed, and documented the 7-week engagement in the internship report.
- Triaged suspicious-login, malware and phishing alerts across 4 SIEM platforms (Splunk, Microsoft Sentinel, QRadar, Elastic) and ran Nessus and OpenVAS vulnerability scans.

PROJECTS

Dispute Desk — LLM Chargeback-Evidence Responder Sep 2026

Razorpay AI Buildathon 2026, Track 02 (AI Risk Manager) — solo submission *Python, LLM APIs*

- Designed an LLM responder that ingests a payment chargeback, assembles evidence from transaction context, and drafts the dispute response an operations team writes by hand.
- Delivered it solo in 5 days with the model inside the decision path rather than generating text at the end of a pipeline; shipped as a public repository with a 5-minute demo.

SetuGuard — Retrieval-Augmented Threat Analysis Pipeline 2026

Grand Finalist, PSB CyberShield 2026 — national finale at IIT Hyderabad, team of 4 *Python, FAISS, Mistral 7B, Flask, XGBoost*

- Engineered a retrieval-augmented pipeline over a 95-package Android corpus: static features drive FAISS retrieval, then a local Mistral 7B emits a schema-constrained JSON report with MITRE mappings and generated YARA rules.
- Eliminated fabricated citations by restricting the schema to a fixed set of valid chunk IDs, plus a grounding gate that caught real hallucinations in live output, then held zero failures across 7 hostile-input cases.
- Pinned generation to temperature 0 and a fixed seed, making every report reproducible run to run.
- Accelerated the dataset-scoring endpoint 11× (7.6 s to 0.66 s) and cut peak memory 2.9× (1.9 GB to 0.66 GB) by lifting 5-fold cross-validation into a precomputed artifact.
- Audited the classifier across 9,082 accounts, isolated label leakage in its graph features, and replaced a single-split score with a 20-seed repeated holdout (median AUCPR 0.271, AUROC 0.872).

PwnBot — Tool-Calling LLM Agent 2026

Python, Groq API

- Implemented an agent that converts a natural-language goal into a sequence of tool calls, parsing raw command-line output into typed objects before the model reasons over it.
- Enforced a scope guard blocking any action outside a declared scope, and split the original script into a multi-module package across 5 refactor phases with keys and model settings externalized.

Encrypted Fraud Scoring Service 2025 – present

Co-developed with a faculty advisor *C++17, Microsoft SEAL, Python, Go, gRPC, Docker, GitHub Actions*

- Developed an inference service scoring fraud risk on bank data the server never decrypts, using CKKS homomorphic encryption; the model reaches 0.979 AUC on the ULB credit-card dataset.
- Compressed per-transaction upload 249× (262 KB to 1 KB) with a Go stream-cipher layer and benchmarked a 4.98× mean speedup over Lattigo at matched ring parameters (N = 8192).
- Released it as an 8-stage container build with digest-pinned images, cosign signatures and a per-release SBOM.

TECHNICAL SKILLS

AI and LLM: Retrieval-augmented generation (FAISS), local model serving (Ollama, Mistral 7B), LLM APIs (Groq, Anthropic), agent tool orchestration, schema-constrained output, output grounding, prompt engineering, evaluation harnesses
Machine Learning: scikit-learn, XGBoost, SHAP, pandas, NumPy; stratified holdout design, AUCPR and AUROC, leakage auditing
Languages: Python, C++17, Go, SQL, Bash
Backend and Infrastructure: FastAPI, Flask, gRPC, PostgreSQL, Docker, GitHub Actions, CI/CD, Linux, REST APIs
Security: OWASP Top 10, static and dynamic analysis, Burp Suite, Nmap, Trivy, Semgrep

LEADERSHIP AND ACHIEVEMENTS

- **Head of Sponsorship, Vivacity**: led the corporate outreach team and closed sponsorships for LNMIIT's annual student festival.
- Authored a survey paper on security testing for machine learning systems. Skilled rank, Hack The Box Academy; access-control bug bounty on Intigriti. Google Cybersecurity Professional Certificate, 2025.