Syed Muhammad Mehdi Abidi

+92-332-3488793 • mehdiali78666@gmail.com • GitHub • LinkedIn • Portfolio • LeetCode

Research Interests

AI safety, responsible AI and trustworthy AI: runtime guardrails for LLM agents, prompt-injection defence, secure tool use and policy enforcement for AI coding assistants, and evaluation of safeguards under adversarial and multilingual inputs.

Education

FAST NUCES (National University of Computer and Emerging Sciences) Karachi, Pakistan
BS in Computer Science, CGPA 3.85/4.0 Jun 2026

- 4x Gold Medalist for a perfect 4.0 SGPA; Rector's Honor List (4x) and Dean's Honor List (3x). Link here

Publications

1. **S. M. M. Abidi et al.**, "SecureSphere-LLM: An Identity-Aware Interception Proxy with Programmable Guardrails for Governing AI Coding Assistants Inside the IDE," *in preparation*, Target venue NAACL 2027 LangCode Workshop.
   - Proposed an identity-aware interception proxy that governs 4 AI coding assistants inside VS Code through 8 programmable guardrails: 31 prompt-injection patterns, 14 secret types, 7 dangerous-code classes and 7 PII entity types.
   - Demonstrated fail-closed enforcement with provider-native refusals and per-request JSONL audit trails, and an evaluation harness over a labelled 78-prompt corpus that blocks 31 of 31 attack samples at a 66 ms median.
2. **S. M. M. Abidi et al.**, "NexusHire: A Candidate-Side Decision-Support Layer over a Five-Stage Job Recommendation Cascade," *in preparation*, Target venue:ACM RecSys 2027.
   - Proposed a five-stage cascade over 123K job postings (Two-Tower retrieval, LambdaMART ranking, MMR diversification, LLM re-ranking) wrapped in candidate-side decision support: counterfactual skill gaps and an independent salary verdict.
   - Demonstrated graceful degradation under cold start, with cold-tier Precision@10 about 35% below warm users rather than collapsing, and a +1 point Precision@5 gain from the reciprocal employer-fit signal.

Research Experience

FAST NUCES Karachi, Pakistan
Final Year Project: SecureSphere-LLM (Research and Development) 2025 - 2026
For visuals and demonstrations See here
AI middleware for secure LLM use inside the IDE, supervised by Dr. Muhammad Rafi

- Built the full stack behind the paper above: a Node.js MITM proxy for VS Code traffic, an NVIDIA NeMo Guardrails service with 8 rails backed by NVIDIA NIM safety models, Presidio-based PII masking, and a React admin dashboard with RBAC, OAuth and JSONL audit trails.

Honours And Awards

- **AgentX Hackathon (Astera Software, Westlake Village, California, USA):** ranked among the top 5 teams. See here
- **Best Final Year Project Nomination (2026):** SecureSphere-LLM nominated for the Best FYP Award at FAST NUCES by an external jury from Turing; to be presented at convocation.
- **International Computer Science Competition (2026):** special honour in the Qualification Round for code that passed all test cases; qualified for the Final Round. See here
- **Data Structures & Algorithms:** perfect score (50/50) on the final examination. See here
- **Upcoming Participation (Nov 2026 - Dec 2026) :** International Machine Learning Challenge (IMLC), CALICO Cup (UC Berkeley), MIT Informatics Cup and NASA Space Apps Challenge (CT Gorton). See here: 1 2 3 4

Teaching Experience

FAST NUCES Karachi, Pakistan
Student Lab Assistant, Data Structures Lab See here Fall 2024

- Taught weekly Data Structures lab sessions to BS(CS) students on linked lists, stacks, queues, trees and graphs; debugged student code live, set lab tasks and graded lab exams.
- Ran visualisation workshops on linked lists, trees and AVL trees, and organised a Data Structures hackathon for students.

Teaching Assistant

I took on these roles, averaging 61+ students per class, to help students grasp difficult concepts, guide them through problems until they could solve them on their own, and give feedback they could act on.

- Computer Networks (CN), Spring 2026: tutorials and office hours on TCP/IP, routing, subnetting and socket programming; graded assignments, quizzes and labs with written feedback. See here
- Theory of Automata (TOA), Fall 2025: weekly problem-solving sessions on DFAs, NFAs, regular expressions, CFGs and Turing machines; designed practice problems and graded with rubrics. See here
- Discrete Structures (DS), Fall 2025: tutorials on logic, proofs, sets, relations, combinatorics and graph theory; prepared worked solutions targeting common mistakes. See here
- Computer Organization & Assembly Language (COAL), Spring 2024: guided students through x86 assembly in MASM (registers, addressing modes, procedures); graded labs and held debugging office hours. See here

Industry Experience

Agento Melbourne, Australia
AI Security Research Lead See here Jun 2026 - Present

Core focus: Guardrails & Runtime Monitoring · Connector-Action Policy Engine · Multi-Provider LLM Routing · Agent Memory & Ontology Design · RAG Knowledge Base

- Selected for the Australian Government-backed Buy Australian AI Partnership Program.
- Guardrails: audited 825 call sites, closed the one unguarded bypass, enforced 3 pre-dispatch invariants in Temporal under CI, and proved coverage by reintroducing the bypass and catching it at the exact call site.
- Policy engine: per-tool allow-lists and RBAC-gated approvals across 10 integrations, enforced at one execution chokepoint.
- LLM routing: complexity-based model selection with semantic caching across OpenAI, Anthropic, Moonshot AI and Gemini; up to 30x lower per-token cost on simple tasks, no loss in output quality.

Imperium Dynamics (Microsoft Partner) Chicago, USA
AI Software Engineer 1 See here Mar 2026 - Jun 2026

- Built scalable LLM applications and LangChain agents on modular Python architectures.
- Designed Microsoft Fabric data pipelines on a medallion architecture, governed end to end with Microsoft Purview for data lineage and cataloguing.
- Delivered AI workflows and cloud automation with Azure AI Foundry and Azure Pipelines, using PDM and secure backend authentication.

Analytic Insider Dearborn, Michigan, USA
Generative Ai Associate Developer See here Feb 2025 - Feb 2026

- Built a Tax Consultant Bot with RAG, LangChain and Pinecone, improving retrieval efficiency by 12%.
- Automated structured data pipelines covering 3,000+ business directories that feed LLM applications.

10Pearls Virginia, United States
Data Science Intern See here Sep 2024 - Dec 2024

- Built an AQI prediction system (Random Forest, Ridge, Logistic Regression) with 78% accuracy, fed by real-time OpenWeather and Hopsworks pipelines.
- Automated CI/CD with GitHub Actions, cutting deployment time by 30%.

Projects

NexusHire: Hybrid Job Recommendation System Python · XGBoost · FAISS · FastAPI · Next.js · Supabase See here

- five-stage cascadeover 123K postings with Two-Tower retrieval, LambdaMART ranking, MMR and Llama 3.1 re-ranking salary model on 244K records nightly ingestion from Adzuna and Greenhouse. .

PathoChat / PathoChat 2.0 LangChain · FAISS · Hugging Face · Mistral-7B · n8n · Pinecone See here

- AI medical chatbot reaching 92% accuracy; extended into an agentic system with n8n workflows and Pinecone vector search. Forked by Dr. Qasem Al-Tashi, Data Scientist at NOV, Houston, USA. See here

Fingerprint Matcher SIFT · PostgreSQL · Kotlin See here

- Full-stack biometric recognition with a Kotlin mobile app; 79% real-time matching accuracy.

Leadership & Activities

IEEE FAST NUCES Student Branch Karachi, Pakistan
President See here 2025 - 2026

- Led the student branch's technical events, 9 workshops and outreach, including IEEE Day 2025 and the IEEEXtreme qualifier.
- IEEE Session Speaker: conducted a hands-on workshop on NVIDIA NeMo Guardrails (input, output and topical rails) for 160+ participants on IEEE Day. See here
- IEEE FAST NUCES x Securiti AI (a Veeam company): organised and hosted "A Gentle Introduction to Backend Engineering", delivered by a Securiti AI engineer to 110+ participants for IEEE Day 2025. See here
- IEEEXtreme Qualifier Round Lead: planned and executed the IEEEXtreme 24-hour programming qualifier at FAST NUCES Karachi, managing 100+ participants (Oct 2025). See here

Google Developer Groups on Campus (formerly GDSC) Karachi, Pakistan
Deputy, Data Science & ML See here 2024 - 2025

- Led 10+ workshops, hackathons and seminars engaging 500+ students.

Certifications

- **Supervised Machine Learning: Regression and Classification**: Stanford University, grade 99.66%. See here
- **Convolutional Neural Networks**: DeepLearning.AI, grade 98.87%. See here
- **Elevating LLM Safety with NVIDIA NeMo Guardrails (Colang 2.0)**: CodeSignal. See here
- **Additional**: IBM (Build AI Agents using MCP) | Microsoft Applied Skills (2) + Copilot Studio Agent Academy | DataCamp (6: GenAI, LLMs, LangChain, Prompt Engineering, ML, ChatGPT). IBM • Badge • MS-1 • MS-2 • Copilot • DataCamp: 1 2 3 4 5 6

Technical Skills and Interests

- **Languages & Backend**: Python, SQL, FastAPI, Django, Flask, Node.js, MongoDB, PostgreSQL
- **ML / AI**: PyTorch, TensorFlow, Keras, scikit-learn, XGBoost
- **GenAI & Agents**: LangChain, RAG, Hugging Face, Pinecone, FAISS, n8n, Temporal, MCP, LLM model routing
- **AI Security & Governance**: NeMo Guardrails, Presidio, prompt-injection defence, PII redaction, Microsoft Purview
- **Cloud & DevOps**: Azure (Blob Storage, AI Foundry, Pipelines), Microsoft Fabric, Docker, Jenkins, GitHub Actions, Git
- **Visualisation**: Matplotlib, Plotly, Seaborn, SHAP, LIME