

Muhammad Talha Yaseen

+92 322 388 4150 | talhayaseen.dev@gmail.com | linkedin.com/in/muhammad-talha-yaseen | Lahore, Pakistan

SUMMARY

Full-stack and AI engineer with a Master's in Computer Science and production experience across backend systems, agentic AI pipelines, and computer vision research. Built and shipped systems using FastAPI, MERN stack, LangGraph, and cloud platforms. Comfortable owning features end-to-end from architecture through deployment, and equally at home in research settings and product-focused engineering teams.

EDUCATION

Lahore University of Management Sciences (LUMS), Lahore May 2026

M.S. Computer Science

Courses: Machine Learning, Computer Vision, Generative AI, Deep Learning, Cloud Computing

Forman Christian College University, Lahore June 2023

B.S. Computer Science

Courses: Machine Learning, Software Engineering, Operating Systems, Compiler Construction

EXPERIENCE

Graduate Research Assistant, CVGL May 2025 - May 2026

LUMS, Lahore, PK

- Developed **RoadWeaver**, a topology-aware road graph extraction system built on SwinV2-Base with a dual-head decoder predicting segmentation masks and node heatmaps in a single forward pass.
- Designed a three-component composite loss (OHem-BCE, Focal Tversky, soft-clDice) and identified a non-obvious prerequisite dependency between loss components, verified independently on both CityScale and SpaceNet datasets.
- Injected two lightweight RGB edge stems (under 1,500 parameters total) at the final two decoder stages to recover fine road boundary detail discarded by 32x encoder compression.
- Achieved **88.1% recall** on CityScale and **89.1% recall** on SpaceNet, compared to 74.1% and 72.8% for the CVPR 2025 state-of-the-art baseline. Paper submitted to **NeurIPS 2026**.

Full Stack Engineer Intern Dec 2025 - Feb 2026

10Pearls Pakistan, Lahore, PK

- Built **Mantiq**, a collaborative MERN stack note-taking platform with rich-text editing and real-time sync; enforced code quality through SonarQube integrated into the CI/CD pipeline.
- Engineered RESTful APIs with Node.js and Express using JWT authentication and MongoDB text indexing; containerized the application with Docker and wrote full unit and integration test coverage with Vitest.
- Collaborated with product and QA teams in agile sprints to ship features from requirements through to production release.

Data Science Intern Sep 2024 - Nov 2024

10Pearls Pakistan, Lahore, PK

- Deployed a **FastAPI** churn-prediction microservice with **83% accuracy**, directly supporting customer retention decisions.
- Built a RAG pipeline using LangChain and Pinecone to power a contextual chatbot over internal documentation.
- Architected and optimized PostgreSQL schemas, improving query performance and maintainability for downstream analytics workflows.

Associate Game Developer Aug 2023 - Aug 2024

MindStorm Studios, Lahore, PK

- Shipped production features for **Hexa Sort**, a 3D mobile puzzle game with **millions of downloads** on Google Play; implemented performance-critical UI components in Unity maintaining smooth frame rates across a wide range of Android devices.
- Designed and ran LiveOps A/B experiments using Firebase Remote Config on economy parameters including coin rewards, booster timing, and difficulty curves; analyzed retention and session length metrics to determine winning variants.
- Contributed to the 3D board interaction system, ensuring responsive touch-to-world raycasting and tile mechanics felt snappy across device tiers.
- Managed Android and iOS build pipelines and integrated Firebase Analytics and Crashlytics for real-time event tracking and crash monitoring.

PROJECTS

Chiron – Multi-Agent Research Validation Platform

Designed and implemented the full agentic backend for an AI platform that validates research hypotheses against existing literature before any resources are committed. Built autonomous deep search agents using LangGraph and Google Gemma APIs capable of traversing research literature at depth. Integrated Tavily for live web retrieval and Pinecone for semantic similarity search over ingested corpora. Orchestrated an adversarial multi-agent review stage where one agent defends the hypothesis and another actively challenges it, surfacing weaknesses before experiment commitment. Built the downstream synthesis pipeline that aggregates agent findings into structured experimental protocols. Stack: FastAPI, LangGraph, LangChain, Pinecone, Tavily, React, TypeScript.

Simulacra – Agentic UX Testing Framework

Building an automated AI-driven UX testing framework that goes beyond functional testing to evaluate cognitive load, UI discov-

erability, and navigational friction. Deploys autonomous Vision-Language Model agents that adopt specific human personas and interact with web interfaces like real users. Implements a Concurrent Think-Aloud protocol where an AI HCI Expert monitors the AI User in real-time, interrogating assumptions before any click is executed. The controller uses Set-of-Mark visual prompting via Playwright for deterministic browser manipulation. Stack: FastAPI, LangGraph, Playwright, Google GenAI SDK.

RAG Chatbot – Medical Insurance

Built an end-to-end document-grounded chatbot using LangChain, Pinecone, and Streamlit. Ingested and processed data from PDFs, DOCX files, and images to support accurate, context-aware query responses over medical insurance documentation.

Qwirl – Social Media Backend

Scalable backend built with FastAPI and PostgreSQL; designed relational schema, implemented RESTful API endpoints with JWT authentication, and optimized data processing pipelines for high-throughput social feed operations.

ACHIEVEMENTS

1st Runner-Up, CeDAR LUMS Startup Weekend 2025 (Powered by Allied Bank). Led a cross-functional team to design and pitch *Committee Aunty*, a blockchain-powered fintech platform digitizing informal savings committees via smart contracts. Recognized for innovation, business feasibility, and leadership among university and industry teams.

TECHNICAL SKILLS

Languages: Python, JavaScript (ES6+), TypeScript, C/C++, C#

AI and Agentic Systems: LangGraph, LangChain, LLM integration, RAG pipelines, multi-agent orchestration, Pinecone, Tavily, Generative AI

ML and Computer Vision: PyTorch, TensorFlow, Scikit-Learn, SwinV2, custom loss design, model training and evaluation

Backend and APIs: FastAPI, Node.js, Express, RESTful APIs, microservices, JWT authentication

Frontend: React, Vite, component-based UI development

Databases: PostgreSQL, MongoDB, Redis, schema design, query optimization

Cloud and DevOps: GCP, Azure (VMs, VNets, CLI), Docker, Kubernetes, GitHub Actions, CI/CD

Tools and Practices: Git, JIRA, SonarQube, Vitest, Playwright, Firebase, agile, code review

CERTIFICATIONS

Machine Learning Specialization - Coursera, Stanford University

Python for Data Science, AI and Machine Learning - Coursera, IBM

Game Development Summer Program - M-Labs