

Hasnain Muavia

Lahore, Pakistan

Phone: +92-345-3241015 [Email: hasnain.muavia2004@gmail.com](mailto:hasnain.muavia2004@gmail.com) [Portfolio: hasnainmuavia.vercel.app](https://hasnainmuavia.vercel.app)
[GitHub: HasnainMuavia1](https://github.com/HasnainMuavia1)

Professional Summary

Lead AI Engineer and Software Engineer with 2+ years of experience architecting and delivering production-grade AI systems, agentic workflows, RAG pipelines, and scalable backend platforms. Proven expertise in Django, FastAPI, LangChain, LangGraph, Docker, AWS EC2, CI/CD pipelines, Redis, Celery, and WebRTC. Led cross-functional teams of 8+ engineers and managed full software development lifecycles from requirements to deployment. Published researcher with peer-reviewed papers in MDPI. Deep experience in medical AI automation, ERP systems, and intelligent business process automation.

Experience

Lead AI Engineer

February 2026 – Present

BuiltPulse

On-site, Lahore

- Promoted from Software Engineer to Lead AI Engineer; lead architecture and delivery of CIA – an enterprise ERP OS with RBAC, HR, Finance (ZohoBooks-equivalent), Project Management (ClickUp-equivalent), WebRTC collaboration, and Compliance modules
- Architect and build end-to-end agentic AI pipelines for medical billing, HR automation, email agents, and social media workflows
- Design multi-agent orchestration systems using LangGraph and LangChain with RAG and production deployment strategies
- Build real-time collaboration systems using WebRTC with presence tracking, online/offline status, and typing indicators (Slack-equivalent)
- Implement distributed asynchronous task processing using Celery, RabbitMQ, and Redis Pub/Sub for high-throughput workflows
- Deploy and maintain services on AWS EC2 with Docker containerization; configure Nginx as reverse proxy and implement CI/CD pipelines for automated deployments
- Monitor system health, performance metrics, and error tracking using Grafana dashboards and Sentry alerting
- Integrate AWS S3 for document storage and Stripe for subscription and billing workflows
- Develop scalable REST APIs using Django REST Framework with role-based access control and multi-tenant architecture

Development Team Lead

March 2026 – Present

Nexxaura | Part-time

United States | Remote

- Lead a cross-functional team of 8 engineers (including senior software engineers) building MedSynthesia – an autonomous AI-powered medical billing CRM that handles eligibility verification, insurance verification, medical coding, denial prediction, claim scrubbing, EOB processing, and self-learning
- Resolve technical bottlenecks, unblock engineers, provide architectural solutions, and maintain code quality standards across the full development lifecycle
- Communicate directly with US-based stakeholders to gather requirements, translate business needs into technical specifications, and deliver sprint outcomes
- Manage task assignment, sprint planning, and team coordination using ClickUp; drive agile ceremonies and process improvement initiatives

Associate AI Engineer

April 2025 – January 2026

Big Binary Tech

Remote

- Developed and deployed AI-powered platforms including POS systems, portals, and invoicing solutions using Django, FastAPI, and Flask
- Designed and implemented email automation agents and Retrieval-Augmented Generation (RAG) systems using LangChain and vector databases
- Delivered end-to-end software features from planning and backend development to testing, deployment, and optimization

Associate Software Engineer

February 2025 – January 2026

Arfa Karim Tech Incubator, DHA Campus

On-site, Lahore

- Built and deployed POS, portal, email agent, AI tutor, and invoicing platforms with scalable Django backend architecture
- Trained industry professionals on AI/ML technologies, LLM integration, and emerging automation tools
- Delivered end-to-end software features from planning and development to testing and production deployment

Gen-AI Engineer

January 2025 – May 2025

Autometa Limited

Hybrid

- Built RAG systems integrating OpenAI and Anthropic models with LangChain and Model Context Protocol (MCP)
- Implemented MCP integrations with Jira, Slack, and GitHub for seamless workflow automation
- Developed AI-powered automation agents for intelligent task automation across business workflows

Technical Skills

AI & Machine Learning: OpenAI, Anthropic, Groq, HuggingFace, LangChain, LangGraph, RAG Systems, Agentic AI, Multi-Agent Orchestration, NLP, Computer Vision, Model Fine-tuning, TensorFlow, PyTorch, Pinecone, LlamaIndex, Autogen, MCP, n8n

Backend & API Development: Python, Django, Django REST Framework, FastAPI, Flask, REST APIs, Celery, RabbitMQ, Redis, Microservices, WebSockets, ReportLab

Frontend & Realtime Systems: JavaScript, React, HTML/CSS, WebRTC, Presence Systems, Real-time Collaboration, Playwright

Cloud, DevOps & Infrastructure: AWS EC2, AWS S3, Docker, Nginx, CI/CD Pipelines, GitHub Actions, Linux, Supabase, Vercel, Stripe Integration

Monitoring & Observability: Grafana, Sentry, Redis Pub/Sub, Performance Monitoring, Error Tracking

Data Engineering & Analytics: SQL, PostgreSQL, MySQL, Pandas, NumPy, Matplotlib, Seaborn, ETL, EDA, Data Wrangling

Tools & Collaboration: Git, GitHub, Jira, ClickUp, Asana, Slack, Cursor, Replit, v0.dev, Bolt, Windsurf

Projects

CIA – Enterprise Resource Planning OS | *Django, LangGraph, WebRTC, Celery, Redis, Confidential | BuiltPulse Docker, AWS EC2/S3, Nginx, GitHub Actions, Sentry, Grafana, CI/CD*

- Architected and led development of a full-scale ERP Enterprise OS with RBAC-based entity management, department hierarchies, HR employee management, and Finance module replicating ZohoBooks functionality (invoicing, expenses, reporting)
- Built Project Management module with full ClickUp-equivalent feature set integrated with AI-powered task suggestions and automation using LangGraph agentic workflows
- Developed real-time collaboration system using WebRTC replicating Slack functionality with channels, direct messaging, presence tracking, online/offline status, and typing indicators
- Implemented Compliance module with document lifecycle management and automatic expiry notifications via Celery beat scheduling and Redis task queuing
- Deployed on AWS EC2 with Docker containerization, Nginx reverse proxy, and CI/CD pipeline for automated build and deployment; used AWS S3 for secure document storage
- Integrated Sentry for real-time error tracking and Grafana for system performance dashboards and observability

MedSynthea – Autonomous AI Medical Billing CRM | *LangGraph, LangChain, Django, Confidential | Nexxaura FastAPI, Celery, Redis, Docker, AWS EC2, GitHub Actions*

- Led end-to-end architecture and team of 8 engineers building an autonomous AI-powered medical billing CRM that self-manages the full revenue cycle
- Designed AI pipeline covering eligibility verification, insurance verification, medical coding, denial prediction, claim scrubbing, EOB processing, and self-learning feedback loops
- Resolved critical technical bottlenecks in AI orchestration and data pipeline architecture, unblocking the team and accelerating delivery timelines
- Communicated directly with US stakeholders to gather requirements, define technical specifications, and deliver sprint milestones; managed all task assignment and sprint execution via ClickUp

PON-P3 Pathogenicity Prediction System | *Django, Celery, RabbitMQ, Pandas, ReportLab*

[GitHub](#)

- Developed bioinformatics web application for Lund University (Sweden) to predict pathogenicity of amino acid variations using asynchronous Django-Celery and RabbitMQ architecture for large-scale batch processing
- Integrated Pandas for genomic data manipulation and ReportLab for automated PDF report generation and email delivery

AI-Generated Cryptocurrency Tweet Detection | *PyTorch, Ollama, Transformers, NLP*

[GitHub](#)

- Developed context-aware transformer model to detect AI-generated cryptocurrency tweets with 99% accuracy; published in MDPI 2025
- Generated training data using local LLMs via Ollama and applied advanced overfitting elimination techniques

Multi-Class Visual Cyberbullying Detection | *TensorFlow, Keras, Computer Vision*

[GitHub](#)

- Built end-to-end deep learning pipeline for multi-class visual cyberbullying detection achieving 98% accuracy; published in MDPI 2024
- Performed full data creation, preprocessing, augmentation, and outlier removal pipeline from scratch

AccountingBot RAG System | *LangChain, Selenium, Pinecone, Vector Database*

[GitHub](#)

- Built intelligent RAG system with Selenium web scraping, vector embeddings, conversational memory, user authentication, and session management using LangChain

EduAI – AI Educational Platform | *Groq, LangChain, Llama Vision, RAG, OCR*

[GitHub](#)

- Comprehensive AI educational platform featuring RAG document Q&A, video lecture interaction, quiz generation with Word export, and OCR using Llama Vision for image/PDF text extraction

Education

University of Management and Technology

Bachelor of Software Engineering – CGPA: 3.73/4.0

October 2021 – July 2025

Lahore, Pakistan

- **Relevant Coursework:** Software Engineering, Web Technologies, Deep Learning, Machine Learning, Data Structures, Algorithms

Superior Group of Colleges

F.Sc Pre-Engineering – Grade: A+

January 2019 – May 2021

Lahore, Pakistan

Certifications

Data Analyst – DataCamp (April 2024)

Associate Data Scientist – DataCamp (April 2024)

Machine Learning Specialization – DeepLearning.AI, Stanford University (July 2023)

Python for Everybody – University of Michigan (July 2023)

Publications

Human or AI? Transformer-Based Detection of AI-Generated Financial Content – MDPI, 2025

Multi-Class Visual Cyberbullying Detection using Deep Neural Networks (CVID Dataset) – MDPI, 2024