

HASSAN ABBAS

AI Engineer — Generative AI · LLM Systems · RAG · Agentic Automation

Karachi, Pakistan

hassanabbas2975@gmail.com

github.com/hassan-abbas-arain

linkedin.com/in/hassan-abbas-arain

Dear Hiring Manager,

I am writing to express my strong interest in AI Engineer roles at your organization. Over the past two years I have shipped production-grade AI systems end to end — from LLM orchestration and retrieval-augmented generation to voice automation and computer vision — first at Teknoloje Solutions and currently at Liquid Technologies, where I am a principal contributor to the Sally meeting-intelligence suite.

At Liquid Technologies, I build the AI backbone of a multi-tenant B2B SaaS platform: governed agent execution where every action is risk-classified and fail-closed with owner approval for high-risk writes; a bidirectional Model Context Protocol integration exposing 44 tenant-scoped tools to Claude Desktop behind an OAuth 2.1 authorization server; a production RAG layer on per-tenant pgvector stores with hybrid retrieval and LLM re-ranking; and a four-provider calendar connector platform that made adding integrations and AI providers dramatically cheaper. Alongside Sally, I shipped Sally Desktop — an offline-first Tauri/Rust meeting recorder with fully on-device Whisper transcription across six GPU backends — plus LeadGenPro, a lead-generation SaaS, and a solo-built read-only Upwork analytics platform.

Previously at Teknoloje Solutions, I delivered five production AI systems in eighteen months: an enterprise RAG platform that cut retrieval time by 60%, a YOLOv8 + ResNet Siamese signature-verification pipeline at 94%+ accuracy with sub-300ms cheque OCR, a multi-camera surveillance system with 95%+ face recognition accuracy, a Twilio and Retell AI voice-agent platform that cut lead response from hours to under two minutes, and a LLaMA-based medical summarisation service turning doctor-patient conversations into structured clinical JSON at sub-3s latency.

What sets my work apart is the combination of AI depth and full-stack ownership. I design the retrieval and agent architecture, implement the NestJS/PostgreSQL backend and the React frontend, and carry systems through deployment, observability, and security hardening — multi-tenant isolation, OAuth 2.1, AES-256-GCM encryption, and append-only audit logs are defaults in my work, not afterthoughts. I hold a BS in Computer Science from FAST NUCES, where my final year project, CureBridge, applied BioBERT and LLM diagnostic assistance to telehealth.

I would welcome the opportunity to discuss how this experience can contribute to your team's AI roadmap. My resume is attached, and my portfolio is live at hassan-abbas.lovable.app. Thank you for your time and consideration.

Sincerely,

Hassan Abbas

AI Engineer · Karachi, Pakistan