

# Ankit Jangwan

https://ankitjang.one/  
LinkedIn: ankit-jangwan

Email : jangwanankitofficial@gmail.com

Mobile : +91-9761370345

GitHub: jangwanAnkit

## EDUCATION

---

### Graphic Era University

*B.Tech in Computer Science*

Dehradun, Uttarakhand

*Jul. 2016 – May 2020*

## SKILLS

---

**AI & ML:** LLM Integration (OpenAI GPT-4, Gemini, Llama), RAG Pipelines (Retrieval-Augmented Generation), Prompt Engineering & Context Management, LangChain, Vector Databases (Qdrant), Embeddings (sentence-transformers, OpenAI), SSE Streaming for Real-Time AI Responses, Function Calling / Structured Outputs, Multi-Provider LLM Architecture with Fallback, AI Cost Optimization & Token Management

**Backend:** Python, Django, FastAPI, Flask, Node.js, Express.js, REST APIs, GraphQL, WebSockets, Async (Celery, Redis), OpenAI API, Gemini API, Stripe, Amazon SES, MSG91, Shopify/WooCommerce APIs, OAuth2/JWT

**Frontend:** React, Next.js, Angular, TypeScript, JavaScript, TailwindCSS, Responsive Design, Component Architecture

**Data & Database:** PostgreSQL, MySQL, Redis (caching & queues), MongoDB, Elasticsearch, ETL Pipelines, Query Optimization, Data Modeling

**Cloud, DevOps & Monitoring:** AWS (EC2, Lambda, SES, RDS, S3, CloudWatch), Azure (VMs, Azure Database, Blob Storage), Docker, CI/CD (GitHub Actions, GitLab CI), Nginx/Gunicorn, SSL/TLS, Infrastructure-as-Code, Datadog, Sentry, Logz.io, cProfile, Django Debug Toolbar

## EXPERIENCE

---

### Software Development Engineer II

*Instahyre*

Remote, India

*Jan 2025 – Jan 2026*

- Full-stack development: Built employer-facing recruitment features using Angular + Django, including candidate screening modules (activities, notes, interview scheduling, evaluations). Architected Angular component structure for unified candidate sidebar view integrating 5+ modules with responsive UI and Django REST API integration.
- Backend performance: Led company-wide Django-to-Jinja migration across 40+ modules; reduced page load from ~1000ms to <200ms (5x improvement). Optimized critical APIs via SQL query rewriting (N+1 elimination, index additions, complex join rewrites), Celery/Redis async offloading, and caching — reduced API response times by 40-60%.
- System design: Architected unified candidate view backend integrating 5+ microservices with efficient data aggregation and low-latency queries (<50ms avg); built Elasticsearch-based resume screening pipeline for candidate matching.
- Mentored SDE-1 engineers on async patterns, DB optimization, and API design; led production incident response using Datadog, Sentry, and Logz.io.

### Business Integration Specialist

*Cloudify*

Remote, India

*Aug 2024 – Dec 2024*

- Full-stack development: Built and deployed Django/Express backends with React/Next.js + TypeScript frontends on AWS Serverless (Lambda, SST) and Azure; managed VMs, load balancers, PostgreSQL RDS/Azure Database instances.
- Delivered production-grade 3rd-party integrations: Amazon SES (retry logic, bounce handling), MSG91 (WhatsApp/SMS, high-volume queue management), Stripe (payments, webhooks, PCI compliance), Celery/Redis (distributed async queues), Shopify/WooCommerce (inventory sync, order webhooks).
- Set up CI/CD pipelines (GitHub Actions, GitLab CI), automated deployments, DB backups, and monitoring (CloudWatch, Azure Monitor); maintained 99.5%+ uptime.

### Lead Software Developer, AI Products

*Saara Inc*

Remote, India

*May 2022 – Aug 2024*

- AI product development: Led SalesGPT — production AI sales assistant for e-commerce merchants built on OpenAI GPT-3.5/GPT-4; owned full lifecycle from product discussions through UI/UX design, backend architecture, and Azure deployment. Built interactive React chatbot UI with real-time messaging and dynamic product recommendations.
- LLM integration: Implemented prompt engineering (system prompts, few-shot examples, context management), rate limiting, and cost optimisation via request throttling, query caching, and batch processing — reduced OpenAI API costs by 40%. Built retry logic with exponential backoff and graceful degradation for production reliability.
- Async AI pipeline: Offloaded LLM inference from the request cycle using Celery/Redis tasks; built ETL pipeline ingesting product catalogs (Shopify/WooCommerce APIs), user sessions, and conversation history into PostgreSQL for fast retrieval in AI context windows.
- Set up custom Datadog metrics for AI response quality (latency, token usage, error rates); A/B tested prompt variants to improve recommendation accuracy.

### Software Engineer — Reporting & Scoring

*SHL*

Remote, India

*Jul 2020 – May 2022*

- Core backend developer on Reporting and Scoring platform (PHP/Symfony-based custom framework): Developed REST APIs, managed databases, and built candidate assessment report generation pipelines. Led final milestone migration of SHL evaluation logic into Aspiring Minds' modules post-acquisition, combining best-of-breed scoring algorithms from both platforms.

## PROJECTS

---

**AI-Powered Portfolio with RAG Assistant** [Personal] | *Python, FastAPI, Qdrant, Gemini API, Cloudflare Workers, GitHub Actions*

- Live production RAG assistant at [ankitjang.one](https://ankitjang.one) — answers questions about my experience using Qdrant vector search, sentence-transformers embeddings, multi-provider LLM fallback (Gemini/OpenAI/Llama), real-time SSE streaming via Cloudflare Workers, and zero-cost inference architecture.

**HealthLab Management Platform** [Personal] | *Go (Gin), PostgreSQL, React, TypeScript, Docker, GORM, Swagger/OpenAPI*

- Pathology lab management system with Go/Gin backend and React admin dashboard. Appointment scheduling, capacity tracking, digital report delivery via Telegram bot.

**Message Scheduler** [Personal] | *Python, Django, Celery, Redis, Amazon SES, Telegram Bot API*

- Automated scheduling system with Django backend supporting email (Amazon SES) and Telegram delivery. Cron-based job execution with Celery/Redis task queues, retry logic, and scalable background processing for recurring high-volume notifications.

**Telegram Chat Manager** [Personal] | *Python, FastAPI, Telethon, Uvicorn, Pydantic, PyInstaller*

- Self-hosted Telegram account manager with web UI and CLI. Bulk delete chats, identify spam accounts, clean up Telegram without sending data to third parties.

**MeetAssist Chrome Extension** [Personal] | *Chrome Extension V3, JavaScript ES6+, Service Workers, Content Scripts, Chrome Alarms API, MutationObserver*

- Chrome extension that auto-joins meetings and switches between Google accounts. Uses MutationObserver to detect meeting times without official API.

**Saara: SalesGPT - AI Sales Assistant** [Saara Inc] | *Python, Django, React, PostgreSQL, OpenAI API, Azure, Celery/Redis*

- Conversational AI assistant for e-commerce with prompt engineering, context management, and Celery-based async LLM inference. Reduced OpenAI API costs by 40% through request throttling, caching, and batch processing. Full Azure deployment with monitoring and auto-scaling.

**Instahyre: Candidate Resume Module** [Instahyre] | *Django, Angular, PostgreSQL, Celery, Redis, Elasticsearch*

- End-to-end candidate screening workspace: activities, notes, stage transitions (review → interview → assessment → offer), evaluator visibility controls, and automated email notifications. Built with Django + Angular + Elasticsearch.

**Cloudify: Lighthouse — Workflow Engine** [Cloudify] | *Python, Django, Celery, Redis, AWS Lambda, REST APIs*

- Zapier-like automation platform connecting HubSpot/Pipedrive CRM to Zoho Invoice, Airtable, Brevo, and Twilio. Distributed async processing with Celery/Redis on AWS Serverless.

**SHL: Reporting & Scoring Platform** [SHL] | *PHP, Symfony, MySQL, REST APIs*

- Enterprise psychometric assessment system generating 10+ report types (proctored, behavioral, aptitude, competency). Led post-acquisition migration of SHL scoring algorithms into Aspiring Minds platform.

**Dynamics Classes** [Personal] | *React, TypeScript, TailwindCSS, Vite, Airtable, Hubspot*

- Educational platform with React/TypeScript. SEO-optimized course pages, Airtable as headless CMS, Hubspot for lead tracking.

For more projects, visit: <https://ankitjang.one/>

## WRITING

---

**Blog:** [ankitjang.one/blog](https://ankitjang.one/blog) — Technical writing on backend performance, AI systems, and engineering

**Case Studies:** [ankitjang.one/case-studies](https://ankitjang.one/case-studies) — Deep dives into production systems I've built

**Dev.to:** @jangwanankit — Articles on Python, Django, and AI engineering

**Medium:** @jangwanankitofficial — Long-form technical articles