

AI-Native Full-Stack + GenAI Engineer

A 24-week, 1-on-1 mentored program that turns beginners into industry-ready engineers.

Pritam Kamble — AI/ML Engineer · pritamkamble.com · +91 72762 79026

Software engineering, AI engineering, and production engineering — taught together, not as separate electives. AI coding tools are your daily driver from week 1, not a topic you get to in week 20. Every week ships something real. JS/TS-first (Next.js + LangGraph) — not a Python/ML research track.

```
// git log --oneline --all
```

The 24-week roadmap

#w01 AI-native developer setup

week 1

Git/terminal, Cursor + Claude Code as your daily driver, prompting for code, AI-output verification mindset — not “AI writes it, I trust it.”

+ first PR merged · unfamiliar repo explained end-to-end

#w02-04 JS→TS deep, DSA I

weeks 2-4

Closures, event loop, async/await under the hood, TypeScript generics + type-safe APIs, problem-solving depth over DSA breadth: hashmaps, two pointers, sliding window, recursion.

+ 30 problems solved · typed API client shipped

#w05-08 Web + frontend engineering

weeks 5-8

HTTP/JWT/cookies/CORS, React hooks, Next.js App Router + server actions, production build: multi-user SaaS dashboard with auth, RBAC, search, pagination.

+ SaaS dashboard shipped · auth + RBAC live

#w09-12 Backend + PostgreSQL

weeks 9-12

Node internals, NestJS (DI, guards, interceptors), PostgreSQL as the primary datastore — relational modeling, indexes, transactions, query plans — plus Redis, queues, rate limiting.

+ production API live · Redis cache + queue

#w13-15 Testing, Docker, cloud

weeks 13-15

Unit/integration/E2E testing — moved ahead of deployment, not bolted on at the end. Docker + GitHub Actions CI/CD. AWS (S3, RDS, Lambda, IAM) deployment architecture, self-deployed.

+ green CI pipeline · self-deployed to AWS

#w16-19 RAG + agents

weeks 16-19

LLM fundamentals (tokens, structured output, cost/latency), RAG with pgvector, advanced RAG (hybrid retrieval, reranking, query rewriting, retrieval evaluation), LangGraph agents — and when not to use one.

+ RAG eval dashboard · agent workflow shipped

#w20-21 AI-native software engineering

weeks 20-21

Supervising AI coding agents on a real repo — task decomposition, code review, tests, security — the skill is reviewing what AI built, not generating it. Plus LLM evaluation, golden datasets, tracing, regression testing.

+ evaluation harness · agent-reviewed PRs

#w22-24 System design + hiring simulation

weeks 22-24

Load balancing, caching, DB scaling, capstone deployed live, then interview-day practice covering every stage real companies actually run: DSA, JS/TS, SQL, backend, React, AI/GenAI, system design, project defense (“show me a bug AI introduced”), HR.

+ capstone live · placement pipeline active

// portfolio recruiters actually check

What you'll have shipped

GitHub repos, live demos, and case-study write-ups — the three things that move you past the first resume screen.

01 Next.js e-commerce store

Full-stack shop with sandbox payments, product catalog, cart, orders, and an admin dashboard.

Next.js 15 · Prisma · Postgres · Webhooks

02 AI codebase analyst

Point it at a GitHub repo and ask “where's auth handled” or “what breaks on payment failure” — get cited answers and an architecture diagram, not another PDF chatbot.

LangGraph · Embeddings · GitHub API · pgvector

03 URL shortener + analytics

Rate-limited API with Redis caching, click analytics, and QR generation.

NestJS · Redis · Postgres

04 DevOps pipeline

Dockerized microservice deployed through GitHub Actions to AWS.

Docker · CI/CD · AWS

Ready to start shipping?

If you commit to the weekly cadence and you're not shipping real, working code by week 3, we sit down and fix the plan together — or you walk, no hard feelings.

Website pritamkamble.com

WhatsApp [+91 72762 79026](https://wa.me/917276279026)

Apply pritamkamble.com/index.html#apply