

AI-Native Full-Stack + GenAI Engineer

Small-group mentorship, self-paced, that turns students 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 the start, not a topic you get to later. Self-paced, not a fixed cohort — each stage ships something real, building toward placement-ready skills, not loyalty to one language or stack.

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

The full roadmap

- #s01 AI-native developer setup** Stage 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
- #s02 JS→TS deep, DSA I** Stage 2
- 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
- #s03 Web + frontend engineering** Stage 3
- 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
- #s04 Backend + PostgreSQL** Stage 4
- 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
- #s05 Testing, Docker, cloud** Stage 5
- 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
- #s06 RAG + agents** Stage 6
- 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
- #s07 AI-native software engineering** Stage 7
- 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

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](#)

// build → test → deploy, applied to your career

The placement pipeline

Stage 1: Milestones

Finish the core curriculum, ship one capstone, clear 60/150 DSA problems, resume v2 done.

Stage 2: Mock interviews

Weekly DSA, project, and HR mocks. Score 7/10+ to unlock referrals.

Stage 3: Referrals & drives

Targeted referrals when you're ready, plus a curated off-campus drive list.

Stage 4: Apply & track

Ten applications a week with an ATS-ready resume, outcomes tracked together.

No one can promise a job — this pipeline maximizes your odds if you consistently hit the milestones. Strong performers get referral opportunities shared directly.