

LIVE COHORT 4

AI ENGINEERING BOOTCAMP

BUILT EXCLUSIVELY FOR
SOFTWARE ENGINEERS



Ship RAG, agents, MCP and voice AI in 12 weekends

✓ For engineers with 2+ years

✓ AI Engineering & DS Bootcamp 4.0 included

✓ 1 year recordings access

✓ Full refund until 1st November

12

WEEKENDS, LIVE

22

LIVE SESSIONS

6

GRADED PROJECTS

1

FINAL SHOWCASE

Starts **Sunday, 25th October 2026** · Saturdays & Sundays, 5 to 8 PM IST

1.5M+ YouTube subscribers

738K+ learners

85K+ paid learners

878 engineers in cohorts 1 to 3

WHY THIS EXISTS

Every team is shipping AI features. **Very few engineers know how to make them hold up.**

Calling a model is easy. Making it retrieve the right document, act safely on GitHub, answer a voice call in under 800ms, pass an eval set and stay inside a budget is the job. This cohort gives you twelve weekends of building exactly that, live, with the people who build it.

WHO IT IS FOR

Software engineers moving into AI roles

Two or more years of writing production code, and a next role that asks for RAG, agents and evals.

Developers adding AI to a product

A feature request with "AI" in it, and no clear way to make it reliable, measurable and cheap enough.

Tech leads and architects

Designing AI systems for a team, and deciding when to fine-tune, when to RAG and how to serve it.

What you need on day one

- ✓ At least 2 years of software engineering or coding experience
- ✓ Comfort writing code. Session 1 includes a fast Python refresh for engineers
- ✓ No machine learning experience needed. We cover AI from an engineering lens
- ✓ A laptop that can run the projects (8 GB RAM minimum, 16 GB recommended)

Who it is not for

- × Freshers or anyone with less than 2 years of experience. Start with the [AI Engineering & Data Science Bootcamp 4.0](#), then come back
 - × Anyone looking for theory without building
- Naming the wrong fit up front saves everyone a refund.*

THE PRODUCTION BAR

You are not writing notebooks. **You are shipping to a production bar.**

Every project comes with a spec a real team would hand you, and you build it the way it would be built at work: tested, traced and ready to show.

THE 12 WEEKENDS

Four phases. 22 live sessions. Six projects to a production bar.

Saturdays and Sundays, 5:00 to 8:00 PM IST. Every session is 3 hours of learning with theory, practical hands-on and a 20-minute break. Every session is recorded.

01

SESSIONS 1 TO 5 · 25TH OCT TO 15TH NOV

Foundations and RAG

Python for AI, LLMs and embeddings, vector databases, LangChain and Docling, then advanced RAG. Diwali weekend off.

02

SESSIONS 6 TO 9 · 21 TO 29TH NOV

Agents, orchestration and evals

Tool calling and the ReAct loop, LangGraph, multi-agent systems, then evals, guardrails and observability.

03

SESSIONS 10 TO 15 · 5 TO 20TH DEC

Production AI

Multimodal AI, deployment on AWS AgentCore, MCP, context and harness engineering, fine-tuning and local models.

04

SESSIONS 16 TO 22 · 26TH DEC TO 16TH JAN

Product, security, voice and inference

AI product thinking, cost and security, system design, voice AI, inference with vLLM, then the final showcase.

THE CALENDAR

S01 SUN 25TH OCT	S02 SAT 31ST OCT	S03 SUN 1ST NOV	OFF 7TH & 8TH NOV	S04 SAT 14TH NOV	S05 SUN 15TH NOV	S06 SAT 21ST NOV	S07 SUN 22ND NOV
S08 SAT 28TH NOV	S09 SUN 29TH NOV	S10 SAT 5TH DEC	S11 SUN 6TH DEC	S12 SAT 12TH DEC	S13 SUN 13TH DEC	S14 SAT 19TH DEC	S15 SUN 20TH DEC
S16 SAT 26TH DEC	S17 SUN 27TH DEC	S18 SAT 2ND JAN	S19 SUN 3RD JAN	S20 SAT 9TH JAN	S21 SUN 10TH JAN	S22 SAT 16TH JAN	

■ Live session, 3 hours □ Diwali weekend, no live sessions ■ Final project showcase

OFFICE HOURS

Stuck between weekends? Mentors run office hours between the live sessions for doubts, project help and code reviews, and the Discord stays active all week. You never have to wait for the next Saturday to get unblocked.

SESSION BY SESSION

From your first LLM app to **agents that act on GitHub**.

01 Phase 1 · Foundations and RAG

25TH OCT TO 15TH NOV

S01 Kickoff and fundamentals · Sun 25th Oct

AI landscape · Python fundamentals · Simple functions · File manipulation · Classes and objects · APIs and decorators · Calling the Groq API

Assignment 0

S03 Vector DBs and RAG · Sun 1st Nov

Introduction to vector databases · Qdrant hands-on · RAG in pure Python · Qdrant advanced operations

S04 LangChain, Docling and chunking strategies · Sat 14th Nov

LangChain intro · RAG in LangChain · Docling document parsing · Hierarchical chunking

S02 LLMs, embeddings and transformer architecture · Sat 31st Oct

How an LLM works · Transformer architecture · Embeddings and semantic similarity · Key parameters in LLMs

OFF Diwali weekend · Sat 7th and Sun 8th Nov

No live sessions. Use it to catch up on the first three sessions.

S05 Advanced RAG techniques, hands-on · Sun 15th Nov

Vector and vectorless RAG architectures · Hybrid RAG · SQL RAG · Graph RAG

Project 1 built in this session

02 Phase 2 · Agents, orchestration and evals

21 TO 29TH NOV

S06 Agentic AI foundations · Sat 21st Nov

What AI agents are · Tool calling and the ReAct loop · Your first agent with LangChain · Routing with Semantic Router · Memory in agents

S08 Multi-agent systems · Sat 28th Nov

Multi-agent systems · Subagents and handoffs · Multi-agent architectures · Subagent design for CI/CD triage and PR review

Project 2 built in this session

S07 Agent orchestration and LangGraph · Sun 22nd Nov

Introduction to LangGraph · Workflows · Conditional branches and loops · A ReAct loop in LangGraph · Hands-on

S09 Evals, guardrails and observability · Sun 29th Nov

Online versus offline evals · Evaluating RAG pipelines · Evaluating agents · Guardrails · How it all comes together

Project 3 built in this session

WHY EVALS COME THIS EARLY

Most AI courses leave evaluation to the end, if at all. Here it lands in week 5, so every system you build after it has to pass an eval set, a guardrail and a trace before it counts.

SESSION BY SESSION · TAKE IT TO PRODUCTION, THEN PRESENT IT

03 Phase 3 · Production AI: multimodal, cloud, MCP and fine-tuning

5 TO 20TH DEC

S10 Multimodal AI · Sat 5th Dec

Architecture of a multimodal model · Extracting data from PDFs and images · A multimodal RAG system

S12 MCP · Sat 12th Dec

What Model Context Protocol is · MCP architecture · Building MCP servers · Building MCP clients

S14 Fine-tuning and local AI · Sat 19th Dec

Fine-tune versus RAG · LoRA and QLoRA training · Synthetic dataset pipelines · Small language models · Ollama

S11 Deployment on AWS · Sun 6th Dec

AWS's agent stack · Deploying LangChain on AgentCore · AgentCore services

S13 Context engineering and harness engineering · Sun 13th Dec

Why context matters in complex AI systems · Context engineering fundamentals · Context propagation in agents · Harness engineering · Harness patterns for tool use, memory and recovery · Hands-on with deepagents

Project 4 built in this session

S15 Capstone project discussion · Sun 20th Dec

Assignment showcase and capstone scoping

04 Phase 4 · Product, security, voice, inference and showcase

26TH DEC TO 16TH JAN

S16 AI product thinking and reinforcement learning · Sat 26th Dec

Turning AI capabilities into valuable products · Thinking like an AI product builder · The intuition behind reinforcement learning · Reward design

S18 Stakeholder management and system design for AI · Sat 2nd Jan

Managing stakeholder expectations · How an AI system is different · AI system design patterns · Production caveats

S20 Inference engineering · Sat 9th Jan

Runtime optimisations: batching, caching, quantization, speculation · Infrastructure: routing, load balancing, autoscaling · Serving an LLM endpoint with vLLM

Project 6 built in this session

S22 Final project showcase · Sat 16th Jan

Present your final project to mentors and peers

Showcase

S17 Advanced AI engineering and security · Sun 27th Dec

Cost optimisation at scale · Semantic caching · Smart model routing · OWASP Top 10 for AI · Defending against adversarial attacks

S19 Voice AI · Sun 3rd Jan

The real-time voice stack · Low-latency voice systems · Interruptions and silence handling · Multi-turn context recovery

Project 5 built in this session

S21 Personal branding and mock interview · Sun 10th Jan

LinkedIn and Twitter positioning · GitHub portfolio storytelling · Live mock interview with Dhaval and Sid

WHAT YOU LEAVE WITH

Six projects that run, and one you present.

PROJECT 1

LegalBot

Advanced RAG over contracts, filings and financial reports, with hybrid search, reranking and role-based access.

PROJECT 2

DevOps agent

A multi-agent system for GitHub: one subagent triages CI/CD, another reviews PRs and acts on them.

PROJECT 3

Eval harness

An offline eval set of 20+ examples, LLM-as-judge and at least one guardrail, added to your RAG or agent.

PROJECT 4

IT helpdesk agent

A compound system: triage and resolver agents, an MCP tool to a ticketing system, RAG over runbooks, persistent memory.

PROJECT 5

Voice or multimodal agent

A voice agent that answers in under 800ms end to end, or one that extracts data from three document types.

PROJECT 6

Inference router

A small model served with vLLM, a router that sends hard queries to a larger model, and a cost report at 10K requests a day.

THE STACK · THE TOOLS AI ENGINEERING TEAMS USE IN 2026

Python

LangChain + LangGraph

Qdrant

Docling

LangSmith

MCP

AWS AgentCore

Ollama + vLLM

Streamlit

Groq

Gemini

Semantic Router

JEV

deepagents

Ragas

NeMo Guardrails

Unsloth

WHAT IS INCLUDED

- ✓ 22 live sessions across 4 phases, 3 hours each
- ✓ 6 graded projects and a final project showcase
- ✓ A live mock interview with Dhaval and Sid
- ✓ 1 year access to all session recordings
- ✓ The full AI Engineering & Data Science Bootcamp 4.0, at no extra cost
- ✓ A private Discord, active all week for doubt clearing
- ✓ Personal branding: LinkedIn, Twitter and GitHub portfolio storytelling
- ✓ Certificate of completion

YOUR INSTRUCTORS

Taught by engineers who build these systems, **not people who describe them.**



Dhaval Patel
FOUNDER, CODEBASICS · EX-NVIDIA

Teaches the core build modules: LLM fundamentals, RAG, multi-agent systems, AI system design and the final project.



Siddhant Pandey
AI RESEARCH ENGINEER, CODEBASICS

Leads the live weekend sessions and practice labs: hands-on coding, project walkthroughs, debugging and real-time problem solving.



Srikanth Dongala
AI MANAGER, ATLIQ

Teaches practical AI engineering and system architecture through real-world projects and hands-on sessions.



Hemanand Vadivel
CO-FOUNDER, CODEBASICS · EX-EDGEWELL

Teaches the soft-skills track: AI product thinking, stakeholder management, LinkedIn strategy and personal branding.

THE SCHEDULE

DAYS	TIME	DATES	BREAK
Saturday & Sunday	5:00 to 8:00 PM IST	25th Oct 2026 to 16th Jan 2027	Diwali, 7th & 8th Nov

Every session is 3 hours of learning with theory, practical hands-on and a 20-minute break. Every session is recorded and uploaded within 24 to 48 hours.

WHAT ENGINEERS IN COHORTS 1 TO 3 SAY

878 engineers have been through it. **Here is what they wrote.**

"Writing the capstone project completely from scratch solidified my learning in a way theory alone never could."

Dinakara Shetty

Senior Engineer / Senior Manager · Cohort 1

"It covers some fairly complex AI engineering concepts, but they are explained in a simplified and practical way. More importantly, there is a strong focus on actually implementing what we learn."

Sunny Dubey

Cohort 2

"There's no spoon-feeding. Instead, they teach the fundamentals and give you the right pointers, so you build a solid foundation and learn to figure things out yourself."

Asish Kumar Mishra

Cohort 2

"What makes this bootcamp different is the strong focus on real-world AI engineering instead of only theory."

Monish Mamilla

Cohort 1

"The program kept me accountable through relevant project assignments, which helped build a real sense of accomplishment."

Manal Kariapper


Cohort 1

"Sid is already trying to cover the new things which are recently published. lot of my queries are cleared now."

Onkar Deshpande


Cohort 3

ON VIDEO · TAP TO WATCH



Sachin Garg

Senior Principal Engineer · Cohort 1



Bindumalini KK

Staff Software Engineer · Cohort 2



Soppadla Ms

Product Manager · Cohort 1



Eswara Prasad

Cohort 2



Vijay Sista

Cohort 2

STRAIGHT FROM THE COHORT CHAT, UNEDITED

Radhika Dhandapani 07:15 PM

R

Thank you @Sid [Team Codebasics] for all your patience and efforts to bring complex concepts in more understandable way. Really appreciate the amount work you putting behind the screens. Thank a lot.

❤️ 2 👍 1 🌟 1 🍷 1

Lavanya 06:31 PM

L

Thanks to the codebasics team. @Sid [Team Codebasics] , @Dhaval Patel (Team Codebasics) . You are a great inspiration!

Kriti K 07:25 PM

KK

thanks to the admin team as well.. the content is nicely organised between discord, recordings and relevant materials to read ..

❤️ 1

Every quote is from a verified review at codebasics.io/testimonials (18 reviews for this bootcamp).

THE PRICE

One price, with a real deadline.

INNER CIRCLE · UNTIL TUESDAY, 20TH OCTOBER

₹42,000

EMI available. Includes the AI Engineering & Data Science Bootcamp 4.0. Already own it? You pay only the difference.

AFTER 20TH OCTOBER

₹54,000

Enrollment closes on Sunday, 1st November.

PRICE BY REGION

REGION	INNER CIRCLE, UNTIL 20TH OCT	FROM 21ST OCT TO 1ST NOV
India	₹42,000	₹54,000
Outside India	US\$660	US\$840

Prices outside India are charged in US dollars at checkout. The AI Engineering & DS Bootcamp 4.0 is included in both.

REFUND

Full refund if you request it by the end of Sunday, 1st November 2026. That is the day after the second live session (Sunday, 25th October and Saturday, 31st October), so you can see how the cohort runs before deciding. Refunds apply to non-EMI enrollments.

QUICK ANSWERS

Do I need ML experience?

No. If you have at least 2 years of software engineering or coding experience, you're ready. We cover AI from an engineering lens.

What if I miss a session?

Every session is recorded and uploaded within 24 to 48 hours, and you keep access for 1 year.

Twelve weekends from now, you could be the engineer who ships it.

[Enroll Now →](#)