# Agentic AI: Build Autonomous Agents

Design, build and evaluate autonomous AI agents with LangChain, private LLMs and DeepSeek.

| DURATION | FORMAT | LEVEL | LABS |
| --- | --- | --- | --- |
| 40 hours | Onsite, online or hybrid | Intermediate | Cloud sandbox included |

## About this programme

Takes engineers from prompting fundamentals to working agents: agent architectures, memory and context, evaluation and safety, then hands-on builds with LangChain, a privately hosted GPT model and DeepSeek for enterprise workflows.

Who it is for: Software engineers, data scientists and solution architects

## What your team will be able to do

- Write and debug prompts for complex, multi-step tasks
- Architect agents with memory, feedback loops and tool access
- Evaluate agents for accuracy, responsiveness, autonomy and safety
- Build LangChain agents that integrate APIs and documents
- Deploy a private LLM and combine it with LangChain for an enterprise workflow

## Before you start

- Working Python and experience calling an LLM API
- Basic understanding of prompts and retrieval (RAG)

## Course outline

### 1. Prompt engineering for agents

- Tokens, temperature and top-p sampling
- Zero-shot, few-shot and chain-of-thought prompting
- Prompt chaining and context-window optimisation
- Evaluating and debugging prompts: relevance, accuracy and bias

### 2. Agentic AI foundations

- What agentic AI is and how it differs from traditional models
- Applications: autonomous agents and workflow automation
- Building blocks: embedding models, retrieval systems and action APIs
- Core components: decision logic, feedback loops and memory

### 3. Prompting and context for agents

- Multi-step decision making and environment awareness
- Dynamic prompting from real-time feedback and multimodal inputs
- Context-window management and efficient memory use
- Lab: designing an agent for a customer support workflow

### 4. Evaluating and scaling agents

- Metrics: accuracy, responsiveness and autonomy
- Ethics, bias and safety risks
- From prototype to deployment
- Lab: a simple agent driven by prompt-based instructions

### 5. Building agents with LangChain

- Chains, agents, memory, prompt templates, retrieval and tools
- Integrating APIs for real-world automation
- Lab: an AI agent for document analysis

### 6. Private LLMs, DeepSeek and enterprise integration

- PrivateGPT: on-premise versus cloud deployment and querying internal data
- DeepSeek's approach to reasoning compared with OpenAI, Llama and Claude models
- Infrastructure, security and compliance for production agents
- Capstone: a PrivateGPT and LangChain agent for an enterprise workflow

## Sample schedule

5 days onsite, or 10 half-day sessions online. We fit days and session lengths to your team's calendar.

| | | |
|---|---|---|
| Day 1 | 8 h | Prompt engineering for agents; Agentic AI foundations |
| Day 2 | 8 h | Agentic AI foundations (continued); Prompting and context for agents |
| Day 3 | 8 h | Prompting and context for agents (continued); Evaluating and scaling agents |
| Day 4 | 8 h | Evaluating and scaling agents (continued); Building agents with LangChain |
| Day 5 | 8 h | Building agents with LangChain (continued); Private LLMs, DeepSeek and enterprise integration |

## Assessment and certification

### How learning is checked

- A short knowledge check on day one and again at the end, so you can see each person's progress
- Graded hands-on labs in every module, reviewed by the trainer
- A final exercise that applies the programme to a problem like your team's own

### Certificate and exams

- Participants who complete the programme receive a Cognitive Networks AI certificate of completion.

## Formats and cohort size

- Onsite: At your office or a venue you choose, in full days.
- Online, live: Instructor-led over video, usually in half-day sessions so people can keep up with work.
- Hybrid: Some days in a room together, the rest online. Useful when a team works across cities.

Cohort size: 15 to 25 people per trainer. Larger teams run as parallel cohorts on the same outline, so everyone learns the same thing.

Pricing: We quote per cohort, based on duration, format and number of cohorts. There is no per-seat price list: tell us your team and we send a proposal within one business day.

### Bring this programme to your team

Request a proposal at cognitivenetworks.ai/contact?programme=agentic-ai

reachus@cognitivenetworks.ai · Bangalore +91 72043 05303 · Dallas +1 (303) 827-7415