



# Generative AI Training Curriculum

From Fundamentals to Production — A Modular Learning Path for Teams  
Adopting AI

10

Courses

5

Learning Paths

3

Skill Levels

# AI Training Curriculum

10 courses across 3 skill levels — Foundation, Applied, and Advanced

10 Courses · 3 Skill Levels · Modular Format

<b>01</b>  <b>Generative AI Fundamentals</b>	<b>FOUNDATION</b>	1 Day
<b>02</b>  <b>Prompt Engineering Essentials</b>	<b>FOUNDATION</b>	1 Day
<b>03</b>  <b>Advanced Prompt Engineering</b>	<b>APPLIED</b>	1 Day
<b>04</b>  <b>AI-Driven Software Development</b>	<b>APPLIED</b>	3-5 Days
<b>05</b>  <b>Working with Claude Cowork</b>	<b>APPLIED</b>	2 Days
<b>06</b>  <b>AI for Product Owners &amp; Project Managers</b>	<b>FOUNDATION</b>	1 Day
<b>07</b>  <b>Building AI Agents: Config &amp; Orchestration</b>	<b>APPLIED</b>	2 Days
<b>08</b>  <b>Building AI Agents: Custom Design &amp; Development</b>	<b>ADVANCED</b>	2-5 Days
<b>09</b>  <b>Architecting LLM &amp; Agentic Applications</b>	<b>ADVANCED</b>	2 Days
<b>10</b>  <b>AI Security, Compliance &amp; Ethics</b>	<b>APPLIED</b>	1-5 Days

# Recommended Learning Paths by Role

Each role follows a tailored sequence — pick the path that matches your team



*Courses can be taken individually or combined into multi-day programmes*



# Generative AI Fundamentals



1 Day · 6–7 Hours



All Roles — Leadership, Business, Technical

An executive-friendly introduction to Generative AI. This course demystifies how modern AI systems work, establishes a shared vocabulary across the organisation, and builds the judgment needed to separate genuine opportunity from hype.

*Prerequisite: None*

## WHAT YOU'LL LEARN

- What Generative AI is, how it differs from traditional ML and rules-based automation, and why it matters now
- How Large Language Models work at a conceptual level — training, tokens, attention, and inference
- Real-world use cases across industries: content generation, summarisation, coding assistance, customer support, and knowledge management
- Limitations and failure modes: hallucinations, data leakage, bias, non-determinism, and knowledge cut-offs
- The GenAI ecosystem: foundation models, APIs, orchestration layers, fine-tuning vs. prompting, and build-vs-buy decisions



# Prompt Engineering Essentials

 1 Day · Hands-On Workshop

 Business Users, Analysts, Consultants, PMs

Learn how to communicate with AI models effectively. This workshop treats prompting as a structured thinking skill — not a collection of tricks — and gives participants a repeatable framework for getting reliable, high-quality outputs.

*Prerequisite: Generative AI Fundamentals or equivalent understanding*

## WHAT YOU'LL LEARN

- How LLMs interpret prompts: tokenisation, context windows, attention, and why word order matters
- The anatomy of a strong prompt: prompt frameworks, instruction, context, examples, constraints, and output format specification
- Zero-shot vs. few-shot prompting: when examples help and when they hurt
- Role prompting, persona assignment, and tone control for domain-specific tasks
- Common failure patterns: vague instructions, prompt injection vulnerabilities, over-constraining, and context window overflow

# </> Advanced Prompt Engineering for Technical audience

 1 Day · Technical Deep Dive

 Engineers, Tech Leads, Developers

Move beyond trial-and-error prompting into systematic, production-grade prompt design. Covers the advanced patterns that make AI output reliable, grounded, testable, and ready for agentic workflows.

*Prerequisite: Prompt Engineering Essentials + basic programming knowledge*

## WHAT YOU'LL LEARN

- **Grounding strategies:** context injection, prompt caching, and RAG prompting: how to structure prompts so the model reasons from real data rather than training assumptions
- **Prompt chaining and multi-step pipelines:** decompose complex tasks into isolated, focused prompts with clean handoffs — the architectural pattern behind every reliable AI workflow
- **Reflexion and self-evaluation:** define explicit rubrics inside the prompt so the model critiques and rewrites its own output against measurable criteria
- **Controlled divergence and convergence:** instruct the model to generate unconstrained options first, then evaluate and filter with a separate prompt — prevents self-censorship and produces more useful output than asking for "creative but practical" in one shot
- **LLM-as-Judge:** run multiple prompt framings of the same problem, then use a second model call to score, compare, and merge outputs — a practical pattern for improving reliability in production pipelines

# </> Advanced Prompt Engineering for Non-technical audience

 Half Day – 1 day · Applied Workshop

 Engineering Managers, Product Managers, Tech Leads

Advanced prompting techniques applied to leadership work — strategic decisions, stakeholder analysis, and high-stakes written output. Same rigour as the Developers course, different context.

*Prerequisite: Prompt Engineering Essentials*

## WHAT YOU'LL LEARN

- **Self-Ask decomposition:** break a complex strategic question into ordered sub-questions, answer each in sequence, and build toward a recommendation the model cannot shortcut in a single response
- **Multi-perspective prompting:** instruct the model to argue from multiple named stakeholder positions — CTO, CFO, customer — then synthesise a balanced recommendation with explicit trade-offs
- **Plan & Refine:** separate planning from generation — outline structure first, generate section by section, then target specific weak sections for revision rather than regenerating everything
- **Reflexion with explicit rubrics:** define the review criteria inside the prompt so self-critique is specific and actionable, not generic agreement with the first draft
- **Meta-prompting:** use the model to design, critique, and improve the prompts themselves — building reusable prompt templates your team can standardize and version



# AI-Driven Software Development



3-5 Days · Workshop + Lab



Software Developers, QA Engineers, Tech Leads

Explore how AI is transforming every phase of the software development lifecycle. This course moves past the hype to provide practical, measured guidance on integrating AI into coding, testing, review, and documentation workflows — with an honest look at where it helps and where it doesn't. Course can be customized for the following tools: **Github Co-Pilot, Claude, Codex, Warp, Cursor.**

*Prerequisite: Basic software development knowledge*

## WHAT YOU'LL LEARN

- AI-assisted coding and refactoring: Copilot, Cursor, Warp, Codex and Claude Code in daily workflows — beyond auto-complete into architectural support. Dive into Spec-Driven Development.
- Test generation, mutation testing, and AI-driven bug detection: improving coverage without inflating test suites
- Automated code review: catching style violations, security issues, and logic errors — and knowing when to override the AI
- Documentation generation: API docs, README files, inline comments, and architecture decision records from code context
- Productivity measurement, intellectual property considerations, and the evolving role of the developer



# Working with Claude Cowork

 2 Days · Workshop + Lab

 Knowledge Workers, Team Leads & Operations — All Roles

Learn to delegate real work to Claude Cowork — Claude's agentic desktop app for knowledge work. This course moves past the hype to show, hands-on, how to hand off multi-step tasks and get finished deliverables back — with an honest look at where it helps and where it doesn't. Day 1 covers using Cowork; Day 2 goes deep on building and deploying. Can be customized by role and industry: Legal, Finance, Small Business, Marketing Ops.

*Prerequisite: None for Day 1 — comfort with everyday work tools. Day 2 builds on Day 1 or equivalent experience.*

## WHAT YOU'LL LEARN

- Delegating effectively: writing tasks that work using Outcome · Sources · Cadence · Format — and knowing when to reach for Cowork instead of chat
- Producing finished deliverables — spreadsheets, documents, and presentations — from raw inputs and your own templates
- Automating recurring work with scheduled tasks that run on their own
- Extending Claude with connectors and plugins (skills, connectors, sub-agents) to make it a specialist for your role
- Working safely: choosing what Claude can access, approval steps, and governance — plus the honest limits
- Building and packaging your own installable plugins — and rolling them out to your team



# AI Opportunities for Product Owners & Project Managers



1 Day · Strategy Workshop



Product Owners, Project Managers, Business Leads

A practical course for the people who decide what gets built and how it gets delivered. Learn how to identify high-impact AI opportunities, write requirements for non-deterministic systems, manage the inherent uncertainty of AI projects, and measure outcomes honestly.

*Prerequisite: Generative AI Fundamentals*

## WHAT YOU'LL LEARN

- Identifying and prioritising AI use cases: the Impact-Feasibility matrix, data readiness assessment, and build-vs-buy analysis. Learn how to use AI-tools and agentic assistance to make your life easier.
- AI vs. automation vs. traditional software: choosing the right approach for each problem and knowing when AI is overkill
- Writing AI-ready requirements: acceptance criteria for probabilistic systems, defining 'good enough', and handling edge cases
- Managing AI projects: experimentation mindset, proof-of-concept scoping, staged rollout strategies, and stakeholder communication
- Measuring value: defining success metrics, baseline measurement, total cost of ownership, and avoiding vanity metrics



# Building AI Agents: Configuration & Orchestration



2 Days · Hands-On Workshop



Consultants, Ops Leads, Tech-Savvy Business Users

Learn how to design and orchestrate AI agents using configuration-driven and low-code tools — without deep software engineering. This course focuses on the design thinking, workflow architecture, and governance needed to deploy agents that are useful, reliable, and safe.

*Prerequisite: Prompt Engineering Essentials*

## WHAT YOU'LL LEARN

- What AI agents are, how they differ from simple chatbots, and when agentic architectures are (and aren't) the right choice
- Agent anatomy: roles, instructions, tools, memory, context management, and handoff protocols
- Multi-agent coordination: workflows, delegation chains, consensus patterns, and human-in-the-loop checkpoints
- Hands-on with no-code and low-code platforms: building practical agents for support triage, research synthesis, and operational workflows
- Risk management: autonomy boundaries, runaway behaviour prevention, cost controls, and output governance



# Building AI Agents: Custom Design & Development

 2-5 Days · Engineering-Focused

 AI Platform Teams, Senior Developers, Innovation Engineers

A deep technical course on designing, building, and operating custom AI agents using agentic platforms, code, frameworks, and APIs. This course goes beyond framework tutorials to focus on the architecture decisions, failure handling, and operational practices that determine whether agents work reliably in production.

*Prerequisite: RAG course + strong software development skills*

## WHAT YOU'LL LEARN

- Agent architectures: ReAct, Plan-and-Execute, reflection loops, and choosing the right pattern for the problem domain
- Tool calling and function execution: schema design, error handling, sandboxing, permission models, and tool composition
- Memory and state management: conversation memory, episodic memory, working memory, and persistent state across sessions
- Multi-agent systems: communication protocols, task decomposition, specialisation strategies, and coordination overhead
- Debugging, testing, and safety: agent tracing and observability, adversarial testing, cost and latency budgets, and kill-switch design



# Architecting LLM & Agentic Applications



2 Days · Architecture Workshop



Mid-to-Senior Software Engineers — New to Agents

Architect a production-grade deep-research agent from first principles — owning the agent loop, tool surface, retrieval strategy, and verification layer yourself. Each layer is motivated by watching the simpler version fail, so you gain the architectural judgment to design agentic systems for your own problems.

*Prerequisite: Comfort with Python and APIs. No prior LLM-engineering experience required.*

## WHAT YOU'LL LEARN

- Agent loop and a disciplined tool surface; query decomposition, planning, and query generation/rewriting
- Retrieval strategy: search-API trade-offs (keyword vs. semantic vs. freshness), reranking, and the live-fetch-vs-pre-crawl decision
- Context engineering and evidence memory under a token budget; single-agent vs. orchestrator-worker multi-agent designs
- Grounding and citation discipline (“no citation, no claim”); anti-hallucination: claim decomposition, NLI checks, runtime judge, self-repair
- Eval-driven development and production hardening: LLM-as-judge, faithfulness metrics, tracing, cost/latency budgets, prompt-injection guardrails

# AI Security, Compliance, and Ethics

 1 Day · Cross-Functional

 All Roles — Legal, IT, Product, Engineering, Leadership

Understand the security risks, legal obligations, ethical responsibilities, and governance requirements that come with deploying AI systems. This course provides actionable frameworks — not theoretical hand-wraving — for responsible AI adoption.

*Prerequisite: Generative AI Fundamentals*

## WHAT YOU'LL LEARN

- AI-specific security threats: prompt injection, data exfiltration, model inversion, adversarial attacks, and supply chain risks
- Data privacy and regulatory compliance: GDPR, AI Act, sector-specific regulations, data processing agreements, and cross-border considerations
- Bias, fairness, and explainability: sources of bias, measurement approaches, mitigation strategies, and transparency requirements
- Responsible AI frameworks: organisational principles, review boards, impact assessments, and incident response procedures
- Governance in practice: acceptable use policies, model cards, audit trails, vendor evaluation criteria, and continuous compliance monitoring

# AI Security for LLM Applications

 2 Day · Technical Deep Dive

 Senior Engineers, AppSec Leads, Architects

You are shipping an application with an LLM at its core. The threat model is different. The transformer has no architectural wall between your system prompt, the user's message, and whatever you retrieved from the database. Every token is processed identically. Defenses must be built around the model, not inside it.

*Prerequisite: Working C# / .NET. No prior ML experience required.*

## WHAT YOU'LL COVER

- Prompt injection and jailbreaking — why input sanitisation alone cannot work, and what architectural boundaries actually contain the blast radius
- Sensitive data in context — how system prompts leak, why retrieval pipelines exfiltrate by design, and how to structure trust zones across the context window
- Insecure output handling — LLM output is untrusted data; how it reaches downstream systems and how to intercept it before it does damage
- Supply chain and model integrity — third-party models, plugins, and fine-tuning pipelines as attack surfaces; what to verify and when
- Detection — what prompt injection, data exfiltration, and abuse look like in your logs, and how to build signals that fire in production

# AI Security for Agentic Applications

 2 Day · Technical Deep Dive

 Senior Engineers, AppSec Leads, Architects

An agent is an LLM with tools, memory, and the ability to act across multiple steps without a human in the loop. Every capability that makes it useful also extends the blast radius. Classical AppSec controls assume code executes under a known identity with a defined permission set. Agents do neither.

*Prerequisite: AI Security for LLM Applications, or equivalent experience.*

## WHAT YOU'LL COVER

- Indirect prompt injection — attackers embed instructions in documents, emails, and web pages the agent will read; how a retrieval step becomes a remote code execution primitive
- Excessive agency and permission scope — why least-privilege is harder when the agent decides what to do next, and how to enforce tool boundaries that survive adversarial inputs
- Memory poisoning — persistent memory as a persistent attack surface; how malicious content survives across sessions and corrupts future decisions
- Multi-agent trust — orchestrator-to-subagent calls carry no verified identity by default; how trust degrades across a pipeline and what signing and scoping looks like in practice
- Human-in-the-loop design — which decisions must pause for approval, how to make checkpoints tamper-evident, and what irreversibility means as a security property