

APPLIED AND GENERATIVE AI

# Introduction to Generative AI

Level: Foundation • 2 days (expandable to 3) • Virtual, In-person

## Overview

Generative AI arrived in most workplaces before anyone explained it. People are now expected to use tools like ChatGPT, Claude, and Copilot without a clear picture of what these systems are, why they sometimes produce brilliance and sometimes confident nonsense, or where they genuinely belong in real work. That gap is dangerous in both directions: it leads some people to trust the tools too much, and others to dismiss them entirely.

This is a hands-on, foundation course. It builds understanding in the order that makes everything else make sense: first what generative AI is and how a large language model actually works, explained in plain language with no mathematics, then an honest account of strengths and limits, and only then hands-on skills, prompting effectively, finding the right uses in your own work, and using these tools responsibly. Rather than survey every product and buzzword, the course goes deep on the core ideas that will still be true when the product names change. Every module includes a lab, and each module builds on the one before it.

## Who Should Attend

- Professionals in any role who want a genuine understanding of generative AI, not just headlines
- Team members whose organizations are rolling out tools like ChatGPT, Claude, or Copilot
- Managers and leads who need to speak knowledgeably about what these tools can and cannot do

## Prerequisites

- None. No technical background, programming experience, or prior AI exposure is required

## What You Will Learn

- Explain in plain language what generative AI is and how it differs from earlier software
- Describe how a large language model works: training, prediction, and context
- State clearly what these systems do well, where they fail, and why they sound confident either way
- Write effective prompts and improve results through iteration
- Identify where generative AI genuinely fits in your own work, and where it does not
- Use generative AI responsibly: privacy, verification, and organizational policy

## Course Outline

### Day one: what it is and how it works

- What Generative AI Is (and Is Not)
  - From software that follows rules to software that learned from examples
  - Generative AI versus the AI that came before it
  - The landscape today: ChatGPT, Claude, Copilot, and image and audio generation
  - Lab: first hands-on session with a leading model, guided by structured tasks

- How Large Language Models Work, in Plain Language
  - Training on text: how a model learns patterns without being told rules
  - Prediction, one word at a time, and why that produces coherent answers
  - The context window: what the model can see when it responds to you
  - Lab: probe a model with experiments that make its inner workings visible
- Strengths and Limits, Honestly
  - What these models are genuinely good at, and often better than expected
  - Hallucination, knowledge cutoffs, and why models never say "I do not know" enough
  - Why the model does not know your business, and what that implies
  - Lab: push a model until it fails, then explain each failure using the day's concepts

### **Day two: putting it to work**

- Prompting: Getting Good Results on Purpose
  - Clear asks: task, context, and the output you want
  - Iterating: treating the first answer as a draft, not a verdict
  - Simple techniques that noticeably improve results
  - Lab: transform weak prompts into strong ones for realistic work tasks
- Where Generative AI Fits in Real Work
  - Common wins by role: writing, summarizing, analyzing, brainstorming
  - Choosing tasks well: what to hand to the model and what to keep human
  - A brief look ahead at agents and where the technology is going
  - Lab: identify and test three concrete uses of generative AI in your own job
- Using It Responsibly
  - What never goes in a prompt: privacy, confidentiality, and sensitive data
  - Verification habits: how to check work you did not do yourself
  - Bias, attribution, and staying inside your organization's policy
  - Lab: review realistic scenarios and decide, with reasons, what responsible use looks like

### **Extended Version**

The three-day version keeps the same gradient and adds breadth and applied practice:

- Hands-on time with image generation and multimodal models
- A fuller introduction to AI agents, as a bridge to *Introduction to Agentic AI*
- Deeper prompting practice, connecting to *Foundations of Prompt Engineering*
- A capstone in which each learner designs, tests, and presents a generative AI workflow for their own role