



Agentic AI

Agentic AI Foundations demystifies the transition from generative AI to autonomous agentic systems. You will explore the technology behind these systems, the risks they introduce, and how to lead effectively in an AI-augmented professional landscape.

Modules	Case Studies	Takeaways	Key Exercises
Module 0 Introduction		<ul style="list-style-type: none"> Understand the consequential transition from generative AI to agentic AI Identify the components of agentic AI 	<ul style="list-style-type: none"> Generate complex, multi-modal outputs with NotebookLM Examine your own AI use
Module 1 The Machine That Can	<ul style="list-style-type: none"> VoiceRun: Voice agents 	<ul style="list-style-type: none"> Develop fluency in the fundamental concepts behind LLMs Recognize the probabilistic nature of LLMs and how it leads to variable outputs Explore the training lifecycle of an LLM 	<ul style="list-style-type: none"> Test how adding context shifts LLM predictions Use NotebookLM to translate written work into visual summaries
Module 2 Agentic AI	<ul style="list-style-type: none"> VoiceRun: Agentic AI and multi-agent systems 	<ul style="list-style-type: none"> Differentiate between “System 1” (fast/instinctive) and “System 2” (slow/deliberate) AI thinking Identify the role of task decomposition and planning in solving complex problems Analyze how multi-agent systems use orchestrator agents to manage specialized sub-agents 	<ul style="list-style-type: none"> Score your own workplace pain points based on iteration, information density, and logic Define the “trigger event,” necessary tool access, and success metrics for a hypothetical agentic system
Module 3 Challenges and Risks of Agentic AI	<ul style="list-style-type: none"> Bias in AI Anthropic’s blackmailing agent 	<ul style="list-style-type: none"> Summarize “The Lethal Trifecta”: the dangerous intersection of private data, untrusted content, and autonomous action Recognize the risk of “brittleness” and how agents fail when faced with novel, out-of-distribution events Understand prompt injection and how malicious instructions can override a model’s safety guardrails Evaluate biases present in AI systems 	<ul style="list-style-type: none"> Practice a five-step verification process on AI outputs Identify the top risk for a proposed agentic system and design a guardrail to prevent it

Modules	Case Studies	Takeaways	Key Exercises
<div style="background-color: black; color: white; padding: 5px; writing-mode: vertical-rl; transform: rotate(180deg);">Module 4</div> <div style="background-color: #d3d3d3; padding: 10px; text-align: center;">Co-Intelligence</div>	<ul style="list-style-type: none"> • VoiceRun • Critical Thinking 	<ul style="list-style-type: none"> • Deconstruct co-intelligence, a partnership where AI executes and humans direct and validate • Explore critical thinking and consider how to maintain critical human skills and expertise • Prepare for a workplace shift from technical execution to strategic orchestration 	<ul style="list-style-type: none"> • Apply a five-step framework to a complex workplace project with AI

Learning requirements: In order to earn a Certificate of Completion from Harvard Online, participants must thoughtfully complete all 5 modules, including satisfactory completion of the associated quizzes, by stated deadlines.