AI
From Ideas to Execution

Course Details

INSTRUCTORS: ENG. AMR HELAL

NO PRIOR EXPERIENCE REQUIRED

E-CERTIFICATE UPON COMPLETION

About The Instructors:

ENG. AMR HELAL

Engineer Amr Hilal is a specialist in artificial intelligence and machine learning, certified by AWS, and holds a diploma in Artificial Intelligence from MIT in the United States. He is the founder of a company specializing in AI and data science training programs.

He has delivered technical training programs for the Ministry of Communications and Information Technology and conducted machine learning training sessions for the Egyptian Cabinet. He has also presented numerous academic lectures at more than 35 Egyptian and international universities, including The American University in Cairo, Cairo University, Ain Shams University, and German University in Cairo.
Additionally, he offers courses on major educational platforms and serves as a lecturer in machine learning at a well-known Egyptian university. He is also the founder of an annual forum dedicated to artificial intelligence enthusiasts at Nile University.

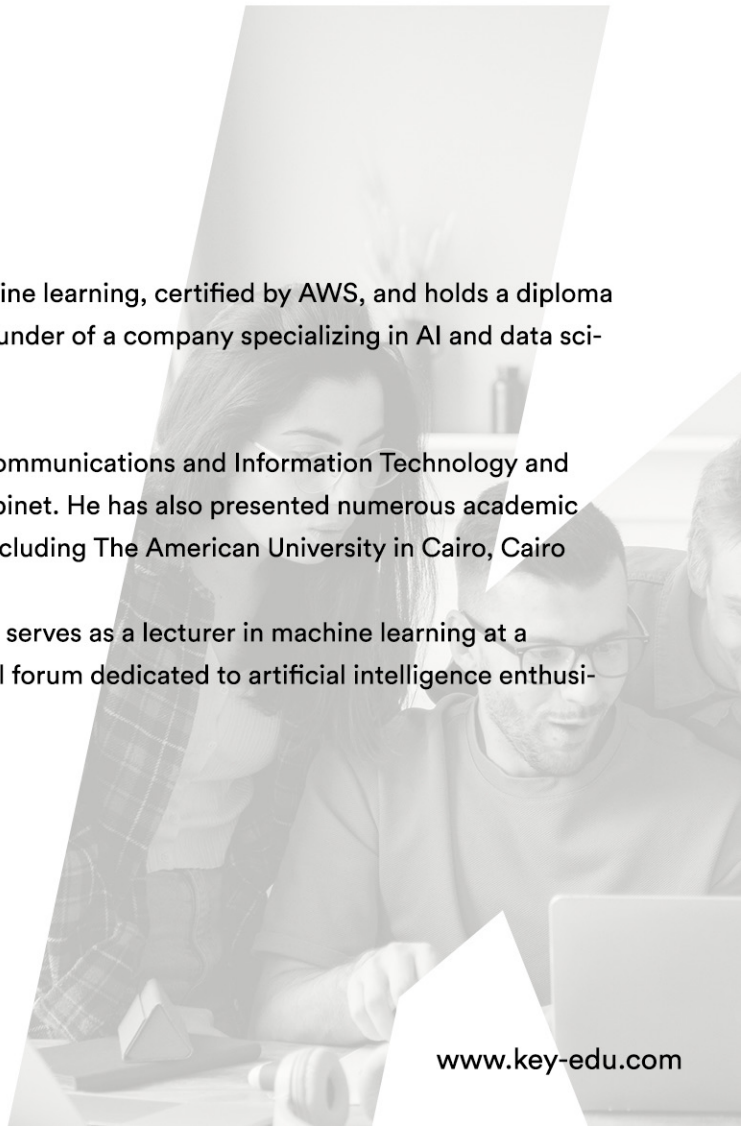

E-Certificate upon completion

MODULE 01Prompt Engineering	***MODULE 02AI Core Tools***	***MODULE 03Extended AI Tools***	***MODULE 04Automation & Agents***	***MODULE 05Advanced Agents***	***MODULE 06AI Strategy***

📌 Course Description

This executive-level AI course is designed to help professionals, managers, entrepreneurs, and C-suite leaders move from understanding AI concepts to implementing real-world AI systems and automation workflows inside organizations.

The course combines strategic AI understanding with practical execution. Participants will learn how to communicate effectively with AI models, use leading AI platforms such as Claude, Gemini, and ChatGPT, build intelligent automations, create AI agents, and develop organization-wide AI adoption strategies.

Across 6 structured modules, learners will explore prompt engineering, AI tools ecosystems, workflow automation, multi-agent systems, AI governance, ROI measurement, and AI roadmap development. The program emphasizes hands-on business applications with real use cases relevant to modern organizations and the MENA region.

By the end of the course, participants will be able to design AI-powered workflows, automate repetitive business operations, evaluate AI opportunities strategically, and lead AI transformation initiatives confidently inside their organizations.

🎯 What You Will Learn

By taking this course, you will learn how to:

🤖 *AI & Automation Skills*

Write powerful prompts that generate accurate and reliable AI outputs

Build structured AI workflows and autonomous AI agents

Automate repetitive business processes using no-code platforms

Create AI-powered decision-making systems

Design multi-agent AI workflows for departments such as HR, Finance, and Sales

Use AI for strategic research, reporting, and executive briefings

🛠️*AI Tools & Platforms*

You will master leading AI tools and ecosystems, including:

Claude

ChatGPT

Gemini

Claude Projects

Claude Code

Apify

Zapier AI

NotebookLM

Perplexity AI

Midjourney

Runway

Make.com

n8n

Prompt & Context Engineering

You will learn:

Prompt engineering fundamentals

Context engineering techniques

Role-based prompting

Few-shot prompting

Chain-of-thought prompting

Tree-of-thought reasoning

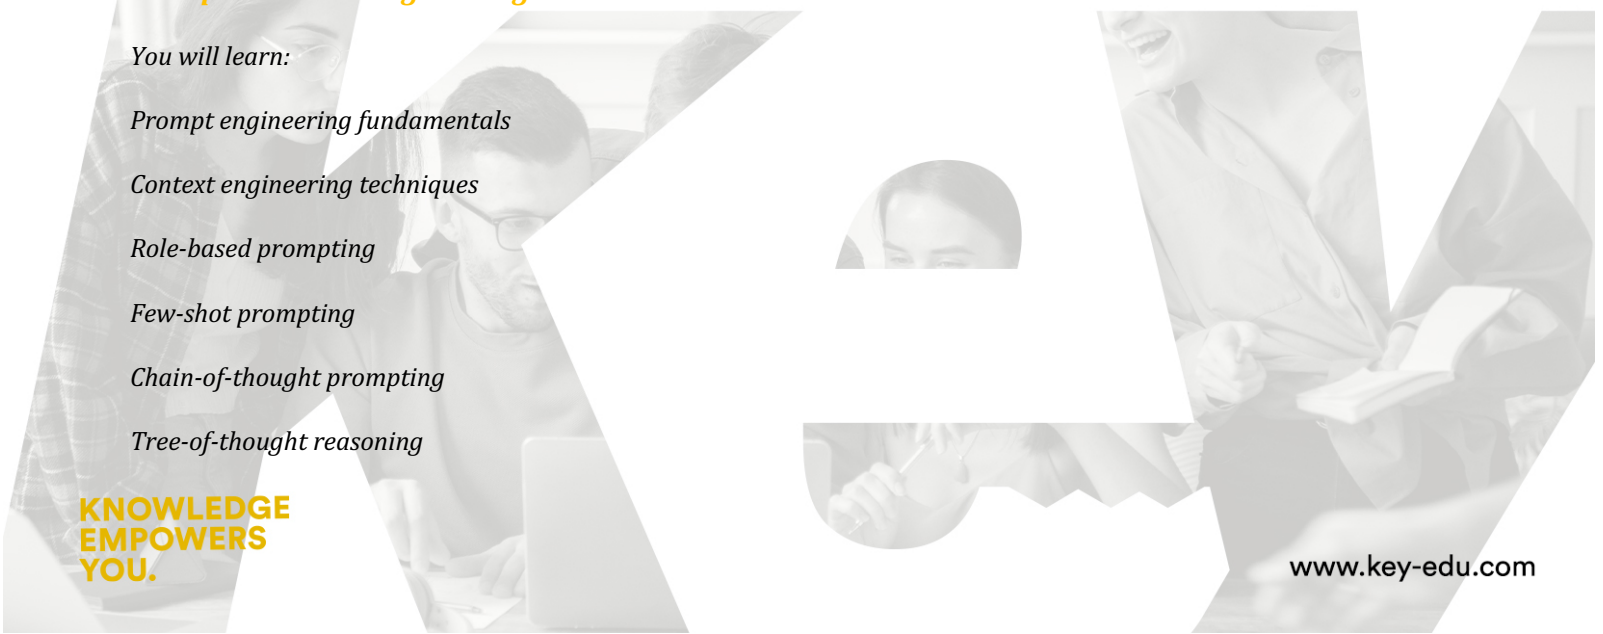

AI constraints and guardrails

Reducing hallucinations with grounded context

📊 *AI Strategy & Governance*

You will understand:

AI ROI measurement frameworks

AI governance principles

Data privacy and compliance

AI risk management

AI adoption roadmaps

AI opportunity audits

Organizational AI transformation strategies

💼 *Business Applications*

Executive intelligence briefings

AI-powered customer support

Automated recruitment workflows

Financial anomaly detection

Lead generation and sales intelligence

Competitive market analysis

AI-enhanced presentations and communication

Automated reporting systems

📝 *Hands-On Practice*

Building AI workflows in Make.com

Designing AI agent systems

Creating automation pipelines

AI-assisted research and synthesis

Building executive dashboards and reports

Working with real business use cases and scenarios

Why You Should Attend This Course

This course is perfect if:

✓*You want to understand AI beyond theory and apply it practically*

✓*You are a manager, entrepreneur, or executive leading digital transformation*

✓*You want to automate workflows and improve operational efficiency*

✓*You want to integrate AI into your business strategy*

✓*You want to save time through intelligent automation*

✓*You want to understand how AI agents and automation systems work*

✓*You want to make better strategic decisions using AI-powered insights*

✓*You want a structured roadmap for AI adoption inside your organization*

By the end of the course, you will have a strong practical and strategic understanding of AI implementation — from writing your first prompt to designing full AI-powered business systems

Who is this course for?

✓*C-suite executives and business leaders looking to lead AI transformation initiatives*

✓*Entrepreneurs and startup founders who want to leverage AI for growth and automation*

✓*Managers and team leaders aiming to improve productivity and operational efficiency*

✓*HR, Finance, Marketing, Sales, and Operations professionals interested in AI-powered workflows*

✓*Consultants and decision-makers seeking data-driven insights and intelligent automation*

✓*Professionals who want to automate repetitive tasks and streamline business operations*

✓*Anyone interested in learning how AI agents, automation systems, and modern AI tools work*

✓*Individuals with little or no technical background who want a practical, business-focused introduction to AI*

✓*Organizations looking to build internal AI capabilities and adopt AI strategically*

Course Content:

This expanded edition of the AI Mastery Program is designed for C-suite leaders who need both strategic literacy and tactical competence in artificial intelligence. Over 16 hours, you will move from foundational prompt engineering through the world's most powerful AI tools — including newly added platforms such as Apify, Zapier AI (Zaps with AI), NotebookLM, Perplexity AI, Midjourney, and Runway — into the architecture of intelligent automation and autonomous AI agents. The program closes with a dedicated module on AI governance, ROI measurement, and building your organization's AI roadmap.

Every session is designed for immediate application. Use cases drawn from MENA business contexts run throughout, and each module culminates in a hands-on scenario you can adapt for your own organization on day one.

16-Hour Schedule at a Glance

Module	Topic	Duration	Cumulative
MODULE 01	Prompt & Context Engineering	3 hrs	0–3 hrs
MODULE 02	AI Tools Ecosystem — Core Platforms	3 hrs	3–6 hrs
MODULE 03	Extended AI Tools (Apify, Zapier AI, Notebook LM, Perplexity, Midjourney, Runway)	3 hrs	6–9 hrs
MODULE 04	Automation & AI Agents (Make.com, n8n)	3 hrs	9–12 hrs
MODULE 05	Advanced Agentic Workflows & Real Use Cases	2 hrs	12–14 hrs
MODULE 06	AI Strategy for C-Suite: ROI, Governance & Roadmap	2 hrs	14–16 hrs
TOTAL		16 hours	

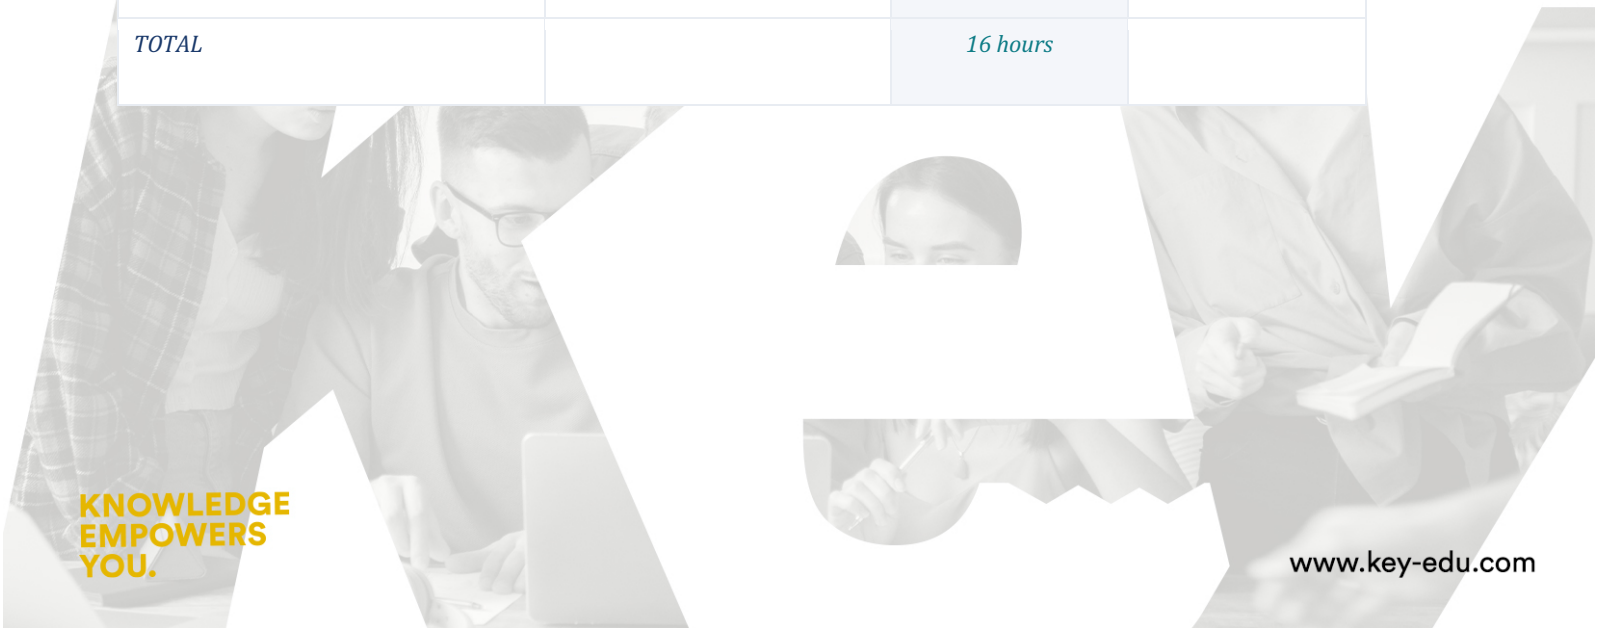

This module establishes the intellectual foundation of the entire course. Before you can build intelligent AI agents, you need to master the language AI speaks — prompts and context. These two disciplines serve distinct and complementary roles in shaping AI behavior.

Part 1: Prompt Engineering

1.1 What Is Prompt Engineering?

Prompt engineering is the practice of designing precise, effective instructions that guide an AI model to produce the exact output you need. For C-suite leaders, it is the skill of delegating to AI the same way you delegate to your best people — clearly, with context, and with defined outcomes.

The specificity of your prompt directly determines the quality, relevance, and reliability of the AI's response.

1.2 Anatomy of a Strong Prompt

Component	What It Does	Example
Role Definition	*Establishes who the AI should behave as*	*Act as a senior financial analyst with 20 years of corporate experience*
Clear Task	*States exactly what you want done*	*Analyze this quarterly report and identify the top 3 risk factors*
Constraints	*Defines what NOT to do or limits scope*	*Do not include data before 2022. Limit to 200 words.*
Output Format	*Specifies how you want the result structured*	*Respond in bullet points with a one-line summary at the top*
Tone & Style	*Shapes the voice and register*	*Use executive-level language. Avoid technical jargon.*

1.3 Constraints and Limitations in Prompts

Constraints are arguably the most powerful — and most underused — element of prompt engineering. They operate on multiple levels:

- *Scope constraints: 'Only analyze data from the MENA region'*
- *Action constraints: 'Do not make recommendations — only report findings'*

- *Format constraints: 'Always respond in a table with three columns'*
- *Behavioral constraints: 'If you are uncertain, say so explicitly. Do not guess.'*
- *Ethical constraints: 'Do not include personal employee data in your response'*

> ***Key Insight: Constraints Become Agent Guardrails***
>
> *Every constraint you engineer into a prompt becomes a rule the AI agent must follow. When you build agents in Module 4, these exact constraints define what your agent can and cannot do autonomously — making prompt engineering not just a communication skill, but your primary tool for AI governance.*

1.4 Advanced Prompting Techniques

Few-Shot Prompting

Provide the AI with 2–3 examples of the input-output pattern you want before asking it to perform the actual task. The AI generalizes from your examples and applies the same structure reliably — ideal for standardized reports, recurring email templates, and classification tasks.

Chain-of-Thought Prompting

Instruct the AI to show its reasoning step by step before arriving at a conclusion. Add 'Think through this step by step' to your prompt. Particularly powerful for financial analysis, risk assessment, and strategic decisions.

Tree-of-Thought Prompting

Ask the AI to explore multiple reasoning branches simultaneously before selecting the best path. Useful for scenario planning and complex decisions with multiple viable options.

Role Stacking

Assign layered roles: 'You are a CFO who also has deep knowledge of MENA tax regulations and speaks to board-level audiences.' The more precise the role stack, the more expert the response.

Part 2: Context Engineering

2.1 What Is Context Engineering?

Context engineering is about everything surrounding your instruction. Context is the full information environment the AI operates within — the documents it can reference, the conversation history it can recall, and the role it has been given at a systemic level.

Your prompt tells the AI what to do; your context tells the AI what it knows. A well-engineered context transforms a general-purpose AI into a domain-expert that understands your business, your terminology, and your goals.

2.2 Layered Context Structure

- *Layer 1 — System Role: Define who the AI is ('You are a senior operations analyst at a retail company in Egypt')*
- *Layer 2 — Background Knowledge: Provide critical business context before the AI answers*
- *Layer 3 — Task-Specific Documents: Upload reports, data, or materials relevant to the current task*
- *Layer 4 — The Prompt: Now give your instruction — the AI is fully equipped*

2.3 Reducing Hallucinations Through Context

When you provide the AI with grounded, real data and explicit instructions like 'Only use the information in the documents I have provided. If you cannot find the answer there, say so,' you constrain the AI to what it actually knows rather than what it might infer.

USE CASE | Executive Data Analysis with Claude
Scenario: You receive a 50-page market research report and need actionable insights for tomorrow's board meeting.
Step 1: Open a Claude Project with your company context already loaded
Step 2: Upload the market research report
Step 3: Prompt: 'You are our Chief Strategy Officer. Analyze this report and extract the top 5 strategic implications for our business in the MENA region. Format as an executive briefing with a one-paragraph summary and five bullet points, each with a recommended action.'
Step 4: Claude returns a board-ready briefing in minutes
Step 5: Follow up: 'Now compare these findings against our Q2 performance data I uploaded earlier' — Claude connects both documents seamlessly

With your prompt and context engineering foundation in place, this module equips you with the practical toolkit. We cover the leading AI platforms — what each does best, how to approach problems with them, and when to reach for which tool.

Part 1: The Claude Ecosystem

2.1 Claude Chat — The Thinking Partner

Claude Chat is your rapid ideation and reasoning engine. Best for: executive briefings, competitive analysis, complex decision framing, legal or policy language review, and high-stakes written communication.

2.2 Claude Projects — Sustained Organizational Intelligence

Claude Projects gives you a persistent workspace where Claude can hold context across multiple sessions. Upload your company documents, set a custom system prompt, and have ongoing conversations that build on each other. Best for: ongoing research projects, maintaining brand or policy consistency, and team-shared AI workspaces.

2.3 Claude Code — Automation Without Developers

Claude Code makes automation accessible to non-technical leaders. Ask Claude Code to process data, build simple scripts, create automated reports, and connect systems — without writing a single line of code yourself. Best for: data processing pipelines, automated report generation, and CSV analysis.

Part 2: Gemini — Deep Research and Synthesis

2.4 What Makes Gemini Unique

Google's Gemini platform has a key capability that sets it apart: real-time web access combined with powerful synthesis. Gemini can actively research the web and synthesize current information from multiple live sources simultaneously — powerful for market intelligence, competitive research, and regulatory monitoring.

2.5 Gemini Deep Research Mode

Deep Research takes a complex research question, breaks it into sub-topics, researches each one across the web, and synthesizes a comprehensive report — all in a single interaction. For executives, this compresses what might take a research team days into minutes.

USE CASE | Competitive Intelligence with Gemini Deep Research
Scenario: You need a comprehensive picture of how your top three competitors are positioning themselves right now.
Step 1: Open Gemini and activate Deep Research mode
Step 2: Prompt: 'Conduct a deep research analysis of [Competitor A], [Competitor B], and [Competitor C]. For each, identify: their current market positioning, recent strategic moves in the past 6 months, pricing signals, key hires, and any product launches.'

Step 3: Gemini researches dozens of sources simultaneously
Step 4: You receive a synthesized, structured competitive intelligence report
Step 5: Export and use directly in your strategy preparation

Part 3: ChatGPT — Speed, Versatility, and Multimodal Power

2.6 Key Strengths of ChatGPT

- *Multimodal input — analyze images, charts, PDFs, and text in the same conversation*
- *Speed and responsiveness for rapid iteration and brainstorming*
- *Strong content generation for presentations, proposals, and communications*
- *GPT-4o voice mode for hands-free executive use cases*
- *Custom GPTs for building specialized, reusable AI tools within your team*

Part 4: Tool Selection Framework

When You Need...	Reach For...	Why
Accurate, nuanced analysis of complex internal documents	***Claude***	*Best reasoning with grounded, safe outputs*
Real-time market, competitor, or regulatory research	***Gemini Deep Research***	*Live web access + multi-source synthesis*
Rapid content creation, brainstorming, multimodal tasks	***ChatGPT***	*Speed, versatility, and creative range*
Small automation scripts and data processing	***Claude Code***	*No developer needed; plain language automation*
Sustained multi-session team projects	***Claude Projects***	*Persistent context and team collaboration*

MODULE 03 · Hours 6–9

Extended AI Tools Ecosystem

Apify · Zapier AI · NotebookLM · Perplexity AI · Midjourney · Runway · Sora

The AI landscape extends well beyond the three platforms in Module 2. This module introduces six powerful specialist tools that, when combined with Claude, Gemini, and ChatGPT, give C-suite leaders a complete operational AI stack. Each tool solves a problem the others cannot.

Part 1: Apify — Intelligent Web Scraping & Lead Intelligence

3.1 What Is Apify?

Apify is the world's largest cloud platform for web scraping and browser automation. It hosts thousands of pre-built 'Actors' (scraping and automation scripts) that collect structured data from websites, social platforms, and online directories — without requiring any coding knowledge. For executives, Apify is a lead intelligence, market data, and competitive monitoring engine.

3.2 Key Business Applications of Apify

- *LinkedIn Lead Scraping: Extract profiles, job titles, company names, and contact signals matching your exact ICP (Ideal Customer Profile) from LinkedIn search results*
- *Company Intelligence: Scrape company websites, press releases, and news for due diligence and competitive monitoring*
- *E-commerce Price Monitoring: Track competitor pricing across platforms automatically on a daily or weekly schedule*
- *Google Maps Lead Generation: Extract business listings, phone numbers, and ratings for B2B outreach in specific geographies*
- *Job Board Monitoring: Track hiring signals from competitors to understand their strategic direction*
- *Social Media Monitoring: Collect mentions, hashtags, and engagement data across platforms*

***USE CASE	LinkedIn Lead Generation with Apify + Claude***
Scenario: Your sales team needs 500 qualified leads matching a specific profile in the MENA region within 48 hours.
Step 1: Open Apify and select the LinkedIn Profile Scraper Actor
Step 2: Configure search parameters: job title, geography (Egypt, UAE, Saudi Arabia), industry, company size

Step 3: Run the Actor — Apify collects structured profile data: name, title, company, LinkedIn URL
Step 4: Export the data as a CSV and open in Claude
Step 5: Prompt Claude: 'Review this list of LinkedIn profiles. Score each contact 1–10 based on fit with our AI training program ICP. Flag the top 50 for immediate outreach and draft a personalized first-line opener for each.'
Step 6: Result: A prioritized, personalized outreach list in under 2 hours

3.3 Apify + Make.com: Automated Intelligence Pipelines

Apify integrates directly with Make.com via webhook or scheduled triggers. You can build fully automated pipelines where Apify collects data daily, Claude analyzes and scores it, and the results are automatically added to your CRM — with zero manual intervention.

Part 2: Zapier AI — Workflow Automation with AI Decision-Making

3.4 What Is Zapier AI?

Zapier is the original no-code automation platform connecting 6,000+ apps. Zapier AI adds intelligent decision-making nodes into these workflows — allowing AI to read inputs, classify information, generate content, and decide next actions within your automated Zaps. Where Make.com offers deep customization, Zapier AI offers breadth of integrations and speed of setup.

3.5 Key Zapier AI Capabilities

- *AI by Zapier: A native AI action node that sends data to an AI model, processes the response, and routes to the next step — without external API configuration*
- *Formatter by Zapier: AI-powered data cleaning and transformation between apps*
- *Zapier Chatbots: Build AI-powered chatbots that connect to your tools (CRM, spreadsheets, databases) in minutes*
- *Document Extraction: AI reads PDFs, emails, and forms to extract structured data automatically*
- *6,000+ Native Integrations: Connect to any business tool your organization uses*

USE CASE | Automated Contract Review Alert with Zapier AI
Scenario: Your legal team receives 30–50 contracts per week. You need AI to flag risk clauses before any human reviews them.
Step 1: Trigger: New PDF email attachment arrives in a dedicated contracts inbox

Step 3: AI by Zapier node: 'Review this contract text. Identify any clauses relating to: unlimited liability, automatic renewal, non-compete restrictions, or governing law outside Egypt. Flag each with a risk level: HIGH / MEDIUM / LOW and a one-sentence explanation.'
Step 4: Zapier routes HIGH risk contracts to a Slack alert for the General Counsel
Step 5: All results are logged to a Google Sheet with timestamp, contract name, and risk summary
Step 6: Result: Legal team only reads contracts flagged HIGH — reducing review time by 70%

Part 3: NotebookLM — Your AI-Powered Knowledge Base

3.6 What Is NotebookLM?

NotebookLM, built by Google, is an AI-powered research and synthesis tool that operates exclusively on the documents you upload. Unlike general AI tools, NotebookLM never goes outside your sources — making it the most accurate and hallucination-resistant tool for working with your own organizational knowledge.

3.7 Key NotebookLM Capabilities

- *Multi-Document Q&A: Upload up to 50 documents and ask questions across all of them simultaneously*
- *Source-Grounded Summaries: Every answer includes clickable citations so you can verify exactly where the AI found the information*
- *Audio Overview ('Podcast Mode'): NotebookLM converts your documents into a spoken-word summary — ideal for absorbing long reports during travel*
- *Mind Maps and Study Guides: Automatically generate visual knowledge maps from complex document sets*
- *Briefing Documents: One-click generation of structured briefings from multi-source inputs*

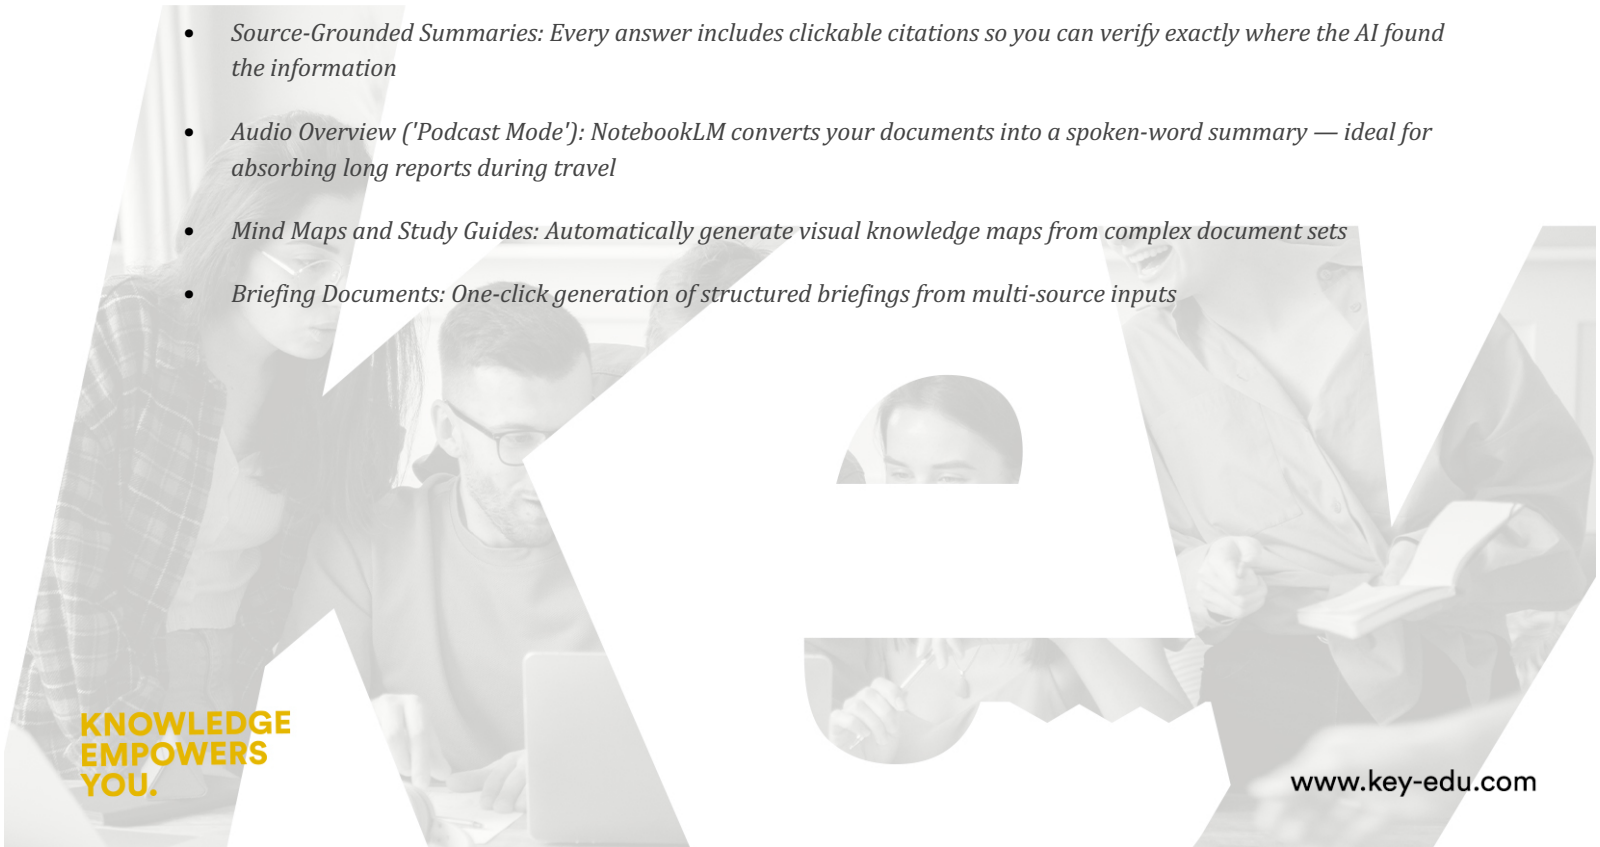

*USE CASE	Board Meeting Preparation with NotebookLM*
Scenario: You have a board meeting in 4 hours and need to synthesize 200 pages of reports, financials, and market data.
Step 1: Upload all documents into a new NotebookLM notebook: annual report, Q3 financials, market research, competitor analysis
Step 2: Ask: 'What are the three most significant risks facing the business based on all these documents?'
Step 3: NotebookLM responds with cited, source-grounded answers — no hallucinations
Step 5: Listen during your commute to the board meeting
Step 6: Result: Complete synthesis of 200 pages in under 30 minutes

Part 4: Perplexity AI — Real-Time Research Engine

3.8 What Is Perplexity?

Perplexity AI is a real-time AI search engine that goes beyond traditional search by providing synthesized, cited answers rather than a list of links. It combines the breadth of web search with the synthesis capability of a language model, making it the fastest tool for grounded, current-events research.

3.9 When to Use Perplexity vs. Gemini Deep Research

Dimension	*Perplexity*	*Gemini Deep Research*
Speed	*Seconds*	*Minutes*
Depth	*Focused answer + sources*	*Comprehensive multi-section report*
Best For	*Quick facts, current news, regulatory updates*	*Deep competitive analysis, strategic research*
Output	*Paragraph with citations*	*Full structured report*

Part 5: Midjourney & Runway — Visual AI for Business Communication

3.10 Midjourney — Executive Visual Communication

Midjourney is the world's leading AI image generation platform. For C-suite executives, it is not a design tool — it is a communication accelerator. Generate presentation visuals, concept illustrations, brand mood boards, and marketing imagery from text descriptions in seconds.

- *Presentation Enhancement: Generate custom, on-brand visuals for board decks and investor presentations*
- *Product Concept Visualization: Quickly visualize new product ideas for stakeholder alignment before development investment*
- *Marketing Campaign Concepting: Generate multiple visual directions for agency briefs in minutes*
- *Brand Mood Boards: Create visual brand references without expensive design sprints*

3.11 Runway — AI Video Generation

Runway is a professional-grade AI video generation and editing platform. For executives, it makes video content accessible without production teams or studios.

- *Gen-3 Alpha: Generate short video clips from text prompts or still images — ideal for social media, internal communications, and product demos*
- *Video-to-Video: Transform existing footage with AI-guided style changes*
- *Automated Corporate Video: Produce training videos, product walkthroughs, and investor update videos at a fraction of traditional production cost*

USE CASE | Corporate Brand Campaign Concepting with Midjourney + Runway
Scenario: Your marketing team needs to present 3 visual campaign directions to the board next week without engaging an agency.
Step 1: In Midjourney, prompt: 'Corporate technology brand campaign visual, MENA market, professional, trustworthy, modern — direction 1: human-centered, warm'
Step 2: Generate 10 images per direction, select the strongest 3 per direction
Step 3: Import selected images into Runway and generate 5-second animated transitions for each
Step 4: Assemble in your presentation tool as animated slides
Step 5: Result: Three fully visualized campaign directions presented to the board — produced in-house in one day

This module is where everything comes together. The prompts you engineered in Module 1, the AI tools you mastered in Modules 2 and 3 — now you will see how to embed them into automated workflows that run on their own. You will build in Make.com and understand n8n as a powerful alternative.

Part 1: Understanding Automation

4.1 The Automation Spectrum

Level	Type	What It Means
Level 1	Rule-Based Automation	Fixed if-then rules. Fast, consistent, brittle with ambiguity.
Level 2	AI-Assisted Automation	AI processes data within a workflow. Handles ambiguity.
Level 3	AI Decision Routing	AI reads input and decides which workflow branch to take.
Level 4	Agentic Workflow	AI agent receives a goal, gathers information, reasons, and takes multiple autonomous actions.

Part 2: Make.com — Your Primary Automation Canvas

4.2 Core Make.com Concepts

- *Scenario: A complete automated workflow — one complete business process*
- *Trigger: The event that starts your scenario (new email, form submission, scheduled time, webhook)*
- *Modules (Nodes): Each node retrieves data, processes it, makes an AI decision, or takes an action*
- *Connections: Authorized links to apps and services (Gmail, Google Sheets, Claude, Slack, WhatsApp, CRMs, Apify)*
- *Routers: Branches in your workflow that direct data to different paths based on conditions*
- *Iterators: Loop through lists of data — process 100 emails, one at a time, automatically*

4.3 Integrating AI Tools into Make.com

Every AI tool covered in Modules 2 and 3 connects to Make.com:

AI Tool	How It Connects in Make.com
Claude (Anthropic)	Native HTTP module with Anthropic API key — send any text, receive AI analysis, summaries, or decisions
ChatGPT (OpenAI)	Native OpenAI module — generate content, classify text, extract data from any input
Gemini (Google)	HTTP module with Google AI Studio API — deep analysis and synthesis tasks
Apify	Native Apify module — trigger scraping Actors, retrieve results, pipe into AI for processing

USE CASE | Customer Service AI Agent in Make.com
Scenario: Build an agent that handles customer service inquiries autonomously — classifying, researching, responding, and escalating intelligently.
Step 1: TRIGGER: New email arrives in the customer service inbox
Step 2: MODULE 1 (AI Classification): Claude reads the email and classifies: type (complaint/inquiry/feedback), urgency (critical/standard/low), customer tier (VIP/standard)
Step 3: MODULE 2 (Route Decision): Make.com routes based on classification — VIP + critical goes to fast track, others to standard flow
Step 4: MODULE 3 (AI Research): Claude searches the internal knowledge base for relevant policies and solution options
Step 5: MODULE 4 (AI Draft Response): Claude drafts a personalized response using gathered context, applying your prompt engineering constraints
Step 6: MODULE 5 (Conditional Action): If urgency = critical, send to manager for approval before sending. If standard, send automatically.

Step 7: MODULE 6 (Logging): Log all details, classification, and response to the CRM automatically
Step 8: RESULT: Your team only touches emails that truly need human judgment — 24/7, consistently

Part 3: n8n — The Open-Source Alternative

4.4 What Is n8n?

n8n (pronounced 'n-eight-n') is an open-source workflow automation platform similar to Make.com but with key differences that make it the preferred choice for organizations with advanced technical needs or data security requirements.

4.5 Make.com vs. n8n — When to Use Which

Dimension	Make.com	n8n
Hosting	*Cloud only*	*Cloud or self-hosted (full data control)*
Data Privacy	*Data passes through Make servers*	*Self-hosted = data never leaves your servers*
Pricing	*Per-operation pricing*	*Open source; hosting cost only (self-hosted)*
Code Flexibility	*No-code / low-code*	*Full JavaScript/Python support in nodes*
Best For	*Fast deployment, wide integrations, non-technical teams*	*Data-sensitive workflows, regulated industries, advanced customization*

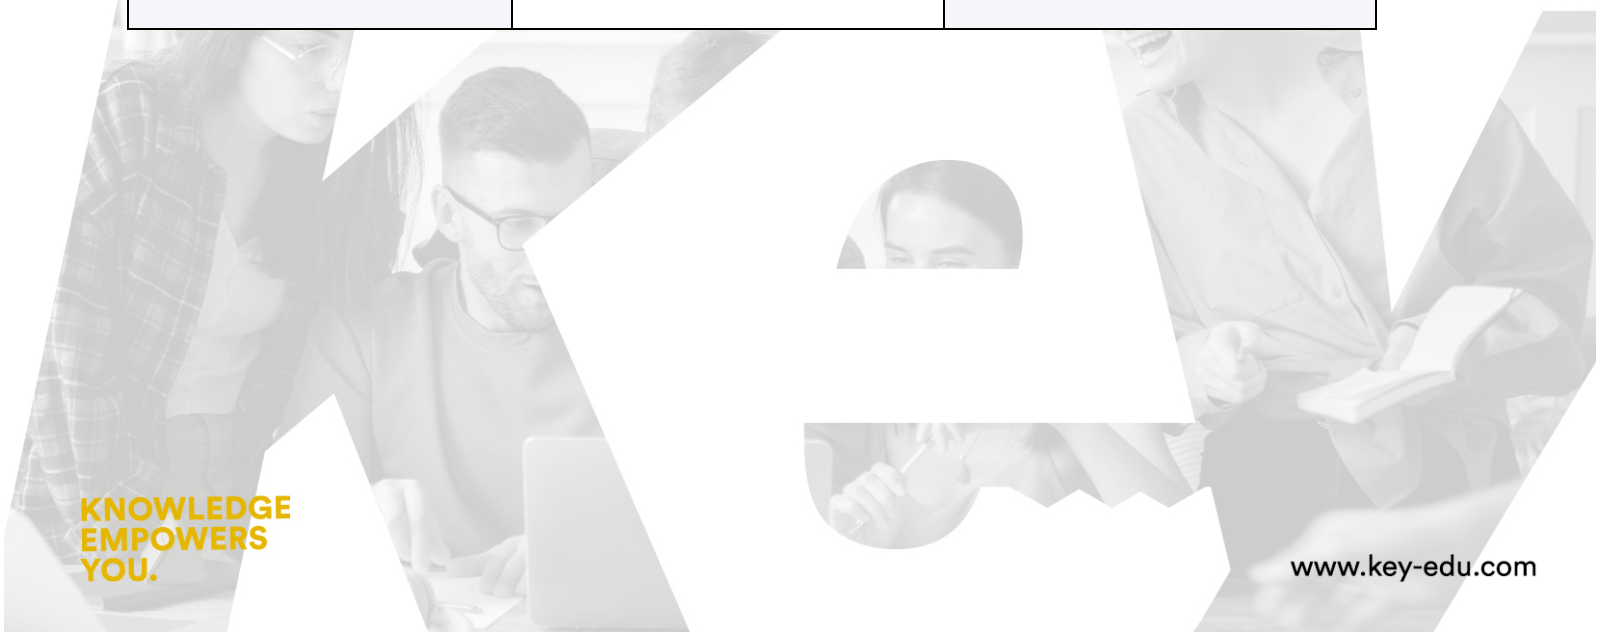

MODULE 05 · Hours 12–14

Advanced Agentic Workflows

Multi-Agent Systems, Memory, Tools, and Real Enterprise Use Cases

You now understand the individual components. This module shows you how to architect sophisticated multi-agent systems — where multiple AI agents collaborate, hand off tasks, and operate as a coordinated intelligence layer across your organization.

Part 1: Anatomy of an AI Agent

5.1 The Four Components of an Agent

Component	What It Does	Example
Perception	Reads and understands incoming information from its environment	Reads incoming emails, Slack messages, form submissions
Memory	Stores information across interactions — short-term (within session) and long-term (across sessions via database)	Remembers past client interactions, preferences, decisions
Reasoning	Uses your engineered prompts and context to decide what action is appropriate given the current state	Decides whether to escalate, respond, research, or wait
Action	Executes the decided action through connected tools and services	Sends email, updates CRM, creates task, posts to Slack

5.2 What Is a Multi-Agent System?

A multi-agent system is a network of specialized AI agents, each with a defined role, that collaborate to complete complex tasks no single agent could handle alone. Think of it as building an AI department — each agent is a specialist, and they hand off work to each other just like a human team.

Multi-Agent Example: Sales Intelligence Department

Agent 1 (Researcher): Receives a company name → scrapes their website, LinkedIn, and news via Apify → passes structured data to Agent 2. Agent 2 (Analyst): Receives structured data → analyzes for strategic fit, decision-maker contacts, and pain points → writes a 3-paragraph company brief → passes to Agent 3. Agent 3 (Outreach

Composer): Receives the brief → drafts a personalized, insight-led outreach email → sends to the sales team for review. The entire pipeline runs automatically when a new prospect is added to the CRM.

Part 2: Real Enterprise Use Cases

5.3 HR Automation Workflow

***USE CASE	AI-Powered Recruitment Pipeline***
Scenario: Automate the full initial screening process for a 50-CV intake for a senior role.
Step 1: TRIGGER: CVs received as email attachments in recruitment inbox
Step 3: Claude scores each CV 1–10 against a role specification you upload as context
Step 4: Scores above 7 → auto-schedule a 15-minute screening call via Calendly integration
Step 5: Scores 4–7 → add to a 'review' list sent to the HR manager weekly
Step 6: Scores below 4 → send a professional decline email automatically
Step 7: All results logged to a Google Sheet with CV attached
Step 8: Result: HR team only meets candidates who have already been AI-screened

5.4 Finance Automation Workflow

***USE CASE	Automated Financial Anomaly Detection & Reporting***
Scenario: Build an agent that monitors your financial data daily and flags anomalies before they become problems.
Step 1: TRIGGER: Scheduled daily at 7:00 AM
Step 2: Agent pulls the previous day's transaction data from your ERP or accounting system via API
Step 3: Claude analyzes the data against your defined thresholds and historical patterns

Step 4: Prompt: 'Identify any transactions that are: more than 30% above the monthly average for that category, duplicate entries, or transactions in unusual currencies or to unfamiliar vendors.'
Step 5: Anomalies flagged → a structured alert is sent to the CFO via WhatsApp and email
Step 6: A daily financial summary is generated and posted to the Finance Slack channel automatically
Step 7: Result: The CFO starts every morning with an AI-generated exception report — problems caught before they compound

5.5 Executive Intelligence Briefing Agent

USE CASE | Automated Morning Executive Brief
Scenario: Every morning at 6:30 AM, you receive a personalized intelligence brief covering your business, industry, and competitors.
Step 1: TRIGGER: Scheduled daily at 6:00 AM
Step 2: Apify scrapes your top 5 competitor websites and news feeds for any updates from the past 24 hours
Step 3: Perplexity API (or Gemini) searches for industry news matching your defined keywords
Step 4: Claude synthesizes: 'Based on the following data, write a 5-point executive brief in the following format: 1. Competitor Watch, 2. Industry Signal, 3. Regulatory Update, 4. Opportunity Flag, 5. Risk Alert.'
Step 5: Brief is sent to your personal email and WhatsApp at 6:30 AM every morning
Step 6: Result: You arrive at every meeting already knowing what changed overnight in your industry

MODULE 06 · Hours 14–16

AI Strategy for C-Suite Leaders

ROI Measurement · AI Governance · Building Your Organizational AI Roadmap

The final module moves from tool mastery to strategic leadership. You leave this course not just knowing how AI works, but knowing how to lead AI adoption inside your organization — measuring ROI, managing risk, and building a roadmap that creates durable competitive advantage.

Part 1: Measuring AI Return on Investment

6.1 The Four ROI Dimensions of AI

ROI Dimension	What to Measure	Example Metric
Time Savings	Hours eliminated from manual tasks	Hours per week saved × team size × average hourly cost
Error Reduction	Reduction in rework, corrections, and escalations	Cost of errors before AI vs. after AI implementation
Speed to Market	Reduction in cycle time for key processes	Days from brief to delivered output — before and after
Revenue Uplift	Incremental revenue from AI-enabled capabilities	Pipeline generated from AI-assisted outreach vs. manual

6.2 The 90-Day AI ROI Framework

Every AI implementation should be measured against a 90-day baseline. Before deploying any AI workflow, document: the current process (who does it, how long it takes, error rate), the cost of the current state (time × cost × volume), and the expected improvement. At 90 days, measure the actual change across all four ROI dimensions and calculate your payback period.

Part 2: AI Governance for Executives

6.3 What Is AI Governance?

AI governance is the set of policies, standards, and controls your organization puts in place to ensure AI is used responsibly, accurately, and in alignment with your values and regulatory obligations. As a C-suite leader, AI governance is your accountability domain.

6.4 The Five Governance Pillars

- *Data Privacy: Define which data can and cannot enter AI systems. No customer PII in external AI tools without explicit policy authorization.*
- *Output Accuracy: Every AI output affecting decisions or communications must have a human review checkpoint in the first 90 days of deployment.*
- *Bias & Fairness: AI systems used in HR, credit, or customer decisions must be audited for systematic bias before deployment.*
- *Vendor Risk: Understand where each AI vendor stores your data and under which legal framework. GDPR / PDPA compliance is non-negotiable.*
- *Change Management: AI adoption fails when people feel replaced. Frame every AI deployment as augmentation — the AI handles the routine; your people handle the judgment.*

Part 3: Building Your AI Roadmap

6.5 The AI Opportunity Audit

Before building your roadmap, conduct an AI Opportunity Audit across your organization. For each department, ask:

- *What tasks take the most time and are the most repetitive?*
- *Where do errors most commonly occur, and what is the cost of those errors?*
- *What decisions are delayed because information is hard to gather or synthesize?*
- *Where do your best people spend time on tasks a capable assistant could handle?*

Each answer is a candidate AI deployment. Rank them by impact (time saved × cost × strategic importance) and feasibility (data availability, tool readiness, team capability). Your top 5 become your 90-day AI roadmap.

6.6 The 4-Phase AI Adoption Roadmap

Phase	Timeline	Focus	Success Metric
1 — Foundation	*Days 1–30*	*Deploy 2–3 simple AI workflows. Build team confidence and literacy.*	*Time saved per week. Zero governance violations.*
2 — Expansion	*Days 31–90*	*Scale to 5–10 workflows. Introduce AI to customer-facing processes with supervision.*	*ROI documented across all deployments. Error rate baseline established.*

4 — Intelligence	*Months 7–12*	*AI becomes a strategic differentiator. Autonomous monitoring, proactive intelligence, decision support.*	*Measurable competitive advantage in speed, cost, and quality vs. peers.*

6.7 Your Personal AI Action Plan — Leave Here With This

By the end of this course, you should be able to answer these five questions for your organization:

1. *Which three processes in my organization would benefit most from AI automation in the next 90 days?*
2. *Which AI tools from this course are the best fit for those three processes?*
3. *What governance policies do I need to put in place before deploying AI to customer-facing or HR processes?*
4. *How will I measure success — and communicate that ROI to my board?*
5. *Who in my organization will I designate as my AI Champion — the person responsible for leading implementation?*

COURSE SUMMARY

You Are Now Equipped to Lead in the Age of AI

16 Hours · 6 Modules · From First Prompt to Full AI Roadmap

Module	Title	What You Can Now Do
01	*Prompt & Context Engineering*	*Speak to AI with precision and give it the institutional knowledge to serve you reliably at scale*
02	*AI Core Tools Ecosystem*	*Select and deploy Claude, Gemini, and ChatGPT strategically for the right problems*
03	*Extended AI Tools*	*Use Apify for lead intelligence, NotebookLM for knowledge synthesis, Perplexity for real-time research, and Midjourney/Runway for visual communication*
04	*Automation & AI Agents*	*Build intelligent automated workflows in Make.com and understand when to use n8n for data-sensitive deployments*
05	*Advanced Agentic Workflows*	*Design multi-agent systems for HR, Finance, Sales, and executive intelligence use cases*
06	*AI Strategy for C-Suite*	*Measure AI ROI, govern AI responsibly, and build your organization's 12-month AI adoption roadmap*

You are now equipped to lead your organization into the age of intelligent automation.