


AI SKILLS


EVERY PROFESSIONAL NEEDS IN 2026



EDITION: 1ST

PUBLISHER: ALKADEMY BOOKS





FREE SAMPLE EDITION

This preview contains the first 1 chapter of the complete book.

BOOK DETAILS

- **Title:** AI Skills Every Professional Needs in 2026
- **Edition:** 1st Edition
- **Publisher:** Alkademy Books
- **Chapters in full book:** 13

CHAPTERS IN THIS PREVIEW

- Chapter 1: Understanding the AI Landscape in 2026

Get the full book at alkademy.com

AI Skills Every Professional Needs in 2026

Alkademy Content Team

1st Edition

Alkademy Books

June 2026

PREFACE

A few years ago, I sat in a meeting that quietly changed the way I thought about my career. A colleague of mine — someone I considered one of the sharpest analysts in our organization — had just spent three weeks building a report that an AI tool completed in forty minutes. She wasn't replaced that day, and nobody said a word about it out loud. But something shifted in the room, and I think we all felt it. The question that followed me home that evening wasn't "Will AI take my job?" It was something more unsettling and, ultimately, more useful: "What do I actually need to know to stay relevant, to stay valuable, and to lead in a world where this technology is no longer coming — it's already here?" This book is my attempt to answer that question honestly, practically, and without the hype that tends to surround every conversation about artificial intelligence.

I wrote this for professionals, not programmers. If you are a marketer, a manager, a consultant, a nurse, a teacher, a lawyer, a financial analyst, or anyone who works in a role that involves thinking, communicating, deciding, or creating — this book was written with you specifically in mind. I have assumed no background in computer science, no prior experience with machine learning, and no familiarity with technical terminology. What I have assumed is that you are curious, that you take your career seriously, and that you are willing to invest genuine effort in building skills that will matter. You do not need to understand how AI works at an engineering level to use it with sophistication and judgment. That distinction is at the heart of everything in these pages.

The philosophy behind this book is simple: fluency over fear, and application over theory. There is no shortage of books that explain what AI is, trace its history, or speculate about its distant future. What I found missing — what I kept searching for and couldn't find — was a clear, grounded guide to what a working professional actually needs to do differently starting now. Every concept introduced here is connected directly to a real professional context. Every skill discussed is one you can begin practicing before you finish the chapter. I believe the best way to understand AI in 2026 is not to study it from a distance but to work alongside it, to develop an instinct for when it helps you and when it misleads you, and to build the judgment that no algorithm can replicate.

The book is organized into four parts. The first part, "Understanding the Landscape," gives you the conceptual foundation you need without drowning you in jargon. You will come away knowing what today's AI tools can genuinely do, where they consistently fail, and why the gap between the two matters enormously in professional settings. The second part, "Core Skills for Every Professional," is the practical center of the book. Here we work through prompt engineering, AI-assisted research and writing, critical evaluation of AI outputs, and the art of integrating these tools into workflows without losing your own voice or judgment. These are the skills that cut across industries and roles, the ones that will serve you whether you work in a hospital, a law firm, a classroom, or a boardroom. The third part, "Advanced Applications by Domain," moves into more specialized territory, with chapters organized around common professional fields. You will find sections relevant to leadership and strategy, communications and content, data and analysis, and client-facing or service roles. You don't need to read all of these — I encourage you to focus on the chapters most relevant to your work and return to the others as your responsibilities evolve. The fourth and final part, "The Human Edge," addresses something I consider the most important and most overlooked dimension of working in an AI-saturated environment: the skills that become more valuable precisely because machines are handling more of the routine. Critical thinking, ethical judgment, creative synthesis, and meaningful collaboration don't diminish in an AI-augmented world — they become the currency that separates professionals who thrive from those who simply survive.

By the time you finish this book, I want you to be able to do several things with confidence. I want you to open any major AI tool and use it purposefully rather than experimentally. I want you to evaluate AI-generated content with a trained eye, catching errors and biases that a less informed colleague might miss. I want you to make thoughtful decisions about when to use AI

and when to trust your own expertise instead. And I want you to walk into any conversation about AI in your organization — whether it's a strategy meeting, a hiring discussion, or a client presentation — feeling equipped rather than anxious.

I should be honest about what this book does not cover, and why. It does not teach you to build AI systems, train models, or write code. Those are valuable skills, but they belong to a different book and a different reader. It also does not attempt to cover every AI tool available — the landscape changes faster than any print publication can track, and naming specific products would date this book before it reached your hands. Instead, I have focused on principles and skills that transfer across tools and will remain relevant as the technology continues to evolve. I have also deliberately set aside the large philosophical debates about AI consciousness, existential risk, and long-term societal transformation. Those conversations matter, but they are not what you need to navigate next Monday morning.

What you need is a place to start — a clear, honest, practical starting point. I hope this book becomes that for you. Let's begin.

To all the lifelong learners, who never stop exploring, questioning, and growing.

To those who dare to teach themselves, and then teach others.

This book is for you.

Contents

Preface	i
1 Understanding the AI Landscape in 2026	1
1.1 Understanding the AI Landscape in 2026	1

CHAPTER 1

UNDERSTANDING THE AI LANDSCAPE IN 2026

1.1 Understanding the AI Landscape in 2026

Imagine arriving at work on a Monday morning to find that your email has already been triaged, your calendar has been reorganized based on project deadlines, a first draft of your weekly report sits in your documents folder, and a brief summary of overnight industry news relevant to your role is waiting in your messaging app. None of this was done by a colleague or an assistant. It was done by a constellation of AI systems working in concert, quietly, while you slept. For millions of professionals, this is not a vision of the future — it is Tuesday. The question is no longer whether artificial intelligence will reshape your professional life. It already has. The only question that matters now is whether you will navigate that transformation with intention and skill, or whether you will find yourself perpetually catching up to a landscape that moved without you.

1.1.1 From Science Fiction to Workplace Infrastructure

There is a particular kind of technological shift that is easy to miss precisely because it happens too gradually and then all at once. Electrification felt like a novelty in 1890 and like plumbing by 1930. The internet felt experimental in 1995 and indispensable by 2005. Artificial intelligence has followed a similar arc, compressed into an even shorter timeframe. What began as a research curiosity confined to university laboratories and well-funded technology companies has, within the span of roughly a decade, become as embedded in professional life as spreadsheets and email.

To understand where we are in 2026, it helps to trace the path that brought us here. The first wave of commercially relevant AI — roughly spanning the 2010s — was dominated by *machine learning* systems that excelled at specific, well-defined tasks. These systems could classify images, detect fraud in financial transactions, recommend products to online shoppers, and filter spam from inboxes. They were powerful within their lanes but brittle outside them. A fraud detection model trained on credit card data could not be repurposed to read medical scans without being rebuilt from scratch. This era produced genuine value, but AI remained largely invisible to most professionals, embedded in backend systems they used without thinking about.

The second wave arrived with the maturation of *large language models* and *generative AI* systems, a period that most observers date to the public release of conversational AI tools beginning in late 2022 and accelerating sharply through 2023 and 2024. Suddenly, AI was not just a background process — it was a collaborator that could write, reason, summarize, translate, code, and converse. Professionals who had never thought about machine learning found themselves using AI tools daily. The technology crossed what researchers call the *adoption threshold*: the point at which a technology becomes easier to use than to avoid.

By 2026, we have entered a third phase, one defined not just by AI that can generate content on demand, but by AI that can pursue multi-step goals autonomously. This is the era of *agentic AI*, and it represents a qualitative leap in capability and complexity. Understanding these three phases — and the distinct types of AI they produced — is the foundational literacy every professional needs to navigate the current landscape intelligently.

Case Study: Consider the transformation of a mid-sized legal services firm over a five-year period. In 2021, the firm used AI primarily in the background: a document management system that automatically categorized filings, and an e-discovery tool that could search large document

sets for relevant terms. Attorneys barely thought of these as AI. By 2023, the firm had adopted a generative AI platform that could draft contract clauses, summarize deposition transcripts, and answer research questions in plain language. Junior associates used it daily. By 2025, the firm had deployed an agentic system capable of receiving a new client intake form and autonomously initiating a sequence of tasks: checking for conflicts of interest across the firm's database, pulling relevant case law, drafting an engagement letter, and scheduling an initial consultation — all before a human attorney reviewed the file. The firm did not replace its attorneys. But it restructured how their time was allocated, and the attorneys who thrived were those who understood how to direct, verify, and build upon the work the AI produced.

1.1.2 The Three Pillars: Narrow AI, Generative AI, and Agentic AI

One of the most common sources of confusion in professional conversations about AI is the tendency to treat it as a single, monolithic technology. In reality, the term “artificial intelligence” describes a broad family of approaches, and the distinctions between them have real consequences for how you use these tools, what you can trust them to do, and where human judgment remains irreplaceable. Three categories deserve particular attention in 2026.

Narrow AI: The Specialist

Narrow AI — sometimes called *weak AI* or *task-specific AI* — refers to systems designed and trained to perform a specific, well-defined function. These systems can be extraordinarily capable within their domain and completely useless outside it. A spam filter is narrow AI. So is the algorithm that determines your credit score, the system that routes your customer service call, the model that detects tumors in radiology images, and the recommendation engine that suggests your next streaming show.

Narrow AI is not glamorous, but it is ubiquitous and consequential. It is responsible for a vast amount of the automated decision-making that shapes professional and personal life. Its defining characteristic is that it optimizes for a specific objective using patterns learned from historical data. This makes it powerful but also brittle in predictable ways: it performs well when new inputs resemble its training data, and it can fail unexpectedly when they do not. A loan approval model trained primarily on data from one demographic can produce discriminatory outcomes when applied more broadly — not because anyone programmed it to discriminate, but because

it learned to replicate patterns in biased historical data.

For professionals, the key skill related to narrow AI is *critical consumption*: understanding what a system was designed to do, what data it was trained on, and where its outputs should be trusted versus scrutinized. This is not a technical skill so much as an analytical habit, and it applies whether you are a nurse reviewing an AI-generated patient risk score, a manager using an AI-driven performance evaluation tool, or a marketer relying on a predictive model to allocate advertising spend.

Generative AI: The Collaborator

Generative AI refers to systems capable of producing new content — text, images, audio, video, code, and more — in response to human prompts. The underlying architecture that powers most modern generative AI systems is the *transformer*, a neural network design introduced in a landmark 2017 research paper that proved extraordinarily well-suited to learning the statistical patterns of language and other sequential data.

What makes generative AI distinctive, and what made its public emergence so disorienting for many professionals, is its apparent fluency. These systems do not retrieve stored answers — they construct responses dynamically, drawing on patterns learned from enormous training datasets. The result can feel remarkably like interacting with a knowledgeable human collaborator. This fluency is both the technology's greatest strength and the source of its most significant failure mode: *hallucination*, the tendency to generate confident-sounding but factually incorrect information.

Example: A financial analyst at an investment firm begins using a generative AI assistant to help draft research summaries. She finds that it dramatically accelerates her workflow — a task that previously took two hours now takes forty minutes. However, she notices that the system occasionally cites statistics or quotes from reports that, when she checks, do not exist. The model has not lied in any intentional sense; it has generated plausible-sounding content that fits the statistical patterns of financial research writing. She adapts her workflow accordingly: she uses the AI to generate structure and draft language, but she verifies every factual claim against primary sources before the report leaves her desk. Her productivity gains are real, but they depend on her maintaining critical oversight rather than delegating judgment entirely.

This example illustrates a principle that runs throughout this book: generative AI is most valuable when it is treated as a powerful but imperfect collaborator, not as an oracle. The professionals

who extract the most value from these tools are those who understand their limitations as clearly as their capabilities.

Agentic AI: The Autonomous Operator

The most significant development in AI capability between 2024 and 2026 has been the emergence of *agentic AI* systems — AI that can pursue multi-step goals autonomously, using tools, making decisions, and adjusting its approach based on intermediate results. Where a generative AI system responds to a single prompt, an agentic system can receive a high-level objective and decompose it into a sequence of actions, executing each in turn.

Agentic systems typically combine a large language model (which handles reasoning and language) with a set of tools (which might include web search, code execution, file management, API calls, and communication platforms) and a planning mechanism (which coordinates the sequence of actions needed to achieve a goal). The result is a system that can, for example, receive the instruction “research our top three competitors’ pricing strategies and prepare a briefing document” and proceed to search the web, read relevant pages, extract pricing information, organize it into a structured format, and produce a polished document — all without further human input.

This capability is genuinely transformative, but it also introduces new categories of risk. An agentic system that misunderstands its objective, encounters an unexpected situation, or makes an error early in a multi-step process can propagate that error through subsequent actions in ways that are difficult to detect and reverse. The professional skill most relevant to agentic AI is not just knowing how to use these systems, but knowing how to *supervise* them: setting clear objectives, defining appropriate boundaries, checking intermediate outputs, and maintaining meaningful human oversight of consequential decisions.

Key Insight

The distinction between generative AI and agentic AI is not merely technical — it is managerial. Generative AI requires you to be a skilled prompt writer and a critical editor. Agentic AI requires you to be a capable delegator and a vigilant supervisor. Both skill sets are valuable; neither is sufficient on its own. As agentic systems become more prevalent in the workplace, the professionals who thrive will be those who can define goals clearly, set appropriate constraints, and evaluate outcomes intelligently — the same skills that make someone an effective manager of human teams.

Dimension	Narrow AI	Generative AI	Agentic AI
Primary Function	Classify, predict, or optimize within a fixed domain	Generate new content from prompts	Pursue multi-step goals autonomously
Typical Interaction	Runs in background; outputs scores or decisions	Single prompt, single response	High-level goal; extended autonomous execution
Key Strength	High accuracy on specific tasks	Fluency, versatility, speed	End-to-end task completion
Key Risk	Bias, brittleness outside training domain	Hallucination, overconfidence	Error propagation, misaligned objectives
Required Professional Skill	Critical consumption, bias awareness	Prompt craft, editorial judgment	Goal setting, supervision, outcome evaluation
Common Examples	Spam filters, credit scoring, medical imaging	ChatGPT, Claude, Gemini, image generators	AutoGPT-style systems, AI copilots with tool access

Table 1.1: Comparison of the three primary AI categories relevant to professionals in 2026.

1.1.3 The Skills Gap and the Opportunity It Creates

Every major technological transition in history has created a period of disequilibrium — a gap between what the technology can do and what the workforce knows how to do with it. The printing press created demand for literacy. The industrial revolution created demand for mechanical and engineering skills. The digital revolution created demand for computer literacy, data analysis,

and eventually software development. AI is creating a new version of this gap, and it is unfolding faster than any previous transition.

The pace of AI capability development has consistently outrun the pace of workforce adaptation. Research from multiple labor market analysts in 2025 found that while AI tool adoption in the workplace had increased dramatically — with a majority of knowledge workers using some form of AI assistance regularly — the depth of that usage remained shallow. Most workers were using AI for simple, low-stakes tasks: drafting emails, summarizing documents, generating first drafts of routine content. Far fewer had developed the deeper skills needed to use AI for complex analysis, to evaluate AI outputs critically, to integrate AI tools into sophisticated workflows, or to understand the strategic implications of AI for their industry and role.

This gap is not a failure of intelligence or motivation. It reflects the natural lag between technological change and educational and organizational response. Companies have been faster to deploy AI tools than to train their employees to use them well. Educational institutions have been slower still to update curricula. The result is a professional landscape in which the ability to use AI effectively has become a genuine differentiator — a skill that commands attention, earns trust, and creates career leverage.

Scenario: Two project managers at the same company are given identical assignments: produce a competitive analysis of five emerging players in their market, with recommendations for strategic response, within one week. The first manager approaches the task the way she always has: she divides the research among her team, holds a synthesis meeting, and writes the final report herself over two days. The second manager uses a combination of AI research tools, generative AI for drafting and structuring, and his own judgment for strategic interpretation and quality control. He produces a more comprehensive analysis in half the time, with richer data visualization and a clearer executive summary. He spends the time he saved stress-testing his recommendations and preparing for the presentation. The difference in output is not primarily a difference in intelligence or experience. It is a difference in AI skill — and that difference is visible to every senior leader in the room.

The professionals who will be best positioned over the next five years are not those who know the most about AI as a technical subject, but those who have developed a practical, applied fluency with AI tools in the context of their own professional domain. This is a learnable skill set, and the chapters that follow are designed to help you build it systematically.

1.1.4 Industries in Transformation: Where AI Is Landing Hardest

AI adoption is not uniform across industries. Some sectors have been early and aggressive adopters; others are in earlier stages of transformation. Understanding where your industry sits on the adoption curve is important context for calibrating your own learning priorities.

Financial services has been among the most aggressive adopters of AI across all three categories. Narrow AI has been embedded in trading, risk management, fraud detection, and credit underwriting for years. Generative AI is now widely used for research synthesis, client communication, regulatory compliance documentation, and financial modeling. Agentic systems are beginning to appear in portfolio management and client advisory contexts. The compliance and regulatory complexity of financial services has created friction but not stopped adoption — it has instead created demand for professionals who understand both AI capabilities and regulatory constraints.

Healthcare presents one of the most consequential AI adoption stories. Narrow AI has demonstrated genuine diagnostic capability in radiology, pathology, and genomics, in some cases matching or exceeding specialist-level performance on specific tasks. Generative AI is transforming clinical documentation, reducing the administrative burden that has been a major driver of physician burnout. Agentic systems are beginning to appear in care coordination and patient communication workflows. The stakes in healthcare mean that human oversight requirements are high and rightly so — but the professionals who understand how to work effectively alongside AI systems are already providing measurably better care with less friction.

Legal services, as noted in the earlier case study, is undergoing rapid transformation. AI is compressing the time required for document review, legal research, contract drafting, and compliance analysis. This is creating pressure on the traditional billing model of large law firms while simultaneously enabling smaller firms and solo practitioners to compete at a higher level than was previously possible.

Education is in the early stages of a transformation that will likely prove as significant as any sector. Generative AI has already disrupted traditional assessment and writing instruction, forcing educators to rethink how they measure learning. More constructively, AI tutoring systems are demonstrating the ability to provide personalized instruction at scale — a capability that has been a holy grail of educational technology for decades.

Manufacturing, logistics, and supply chain have long been AI adopters at the operational level,

using narrow AI for quality control, demand forecasting, and routing optimization. Generative and agentic AI are now entering these sectors at the planning and management layer, enabling more sophisticated scenario analysis and adaptive decision-making.

The common thread across all of these sectors is that AI is not replacing entire professions so much as it is restructuring the work within them — shifting time and attention from lower-order tasks toward higher-order judgment, relationship management, and strategic thinking. The professionals most at risk are not those in complex, judgment-intensive roles, but those whose primary value has been the execution of routine, well-defined cognitive tasks that AI can now perform faster and cheaper.

1.1.5 Building the Habit of Continuous AI Literacy

One of the most consequential mistakes professionals make in response to technological change is treating adaptation as a one-time event. They attend a workshop, read a book, complete a course, and consider themselves updated. This approach may have been adequate for slower-moving technological transitions. It is entirely inadequate for AI in 2026.

The pace of development in AI is genuinely unprecedented. Capabilities that were considered years away are arriving in months. Tools that represented the state of the art in early 2024 have been superseded multiple times. New categories of application emerge regularly. This is not a reason for despair — it is a reason to build different habits. Rather than seeking a single moment of comprehensive understanding, the goal is to establish a sustainable rhythm of ongoing engagement with AI developments that keeps you meaningfully current without consuming your entire professional attention.

Effective continuous AI literacy has several components. The first is *curated information intake*: identifying a small set of high-quality sources that cover AI developments relevant to your domain and engaging with them regularly. The goal is not to read everything, but to read the right things consistently. For most professionals, this means two or three newsletters or publications that filter signal from noise, supplemented by occasional deeper dives into topics that are directly relevant to their work.

The second component is *active experimentation*. Reading about AI capabilities is far less valuable than using AI tools directly. The professionals who develop the deepest fluency are those who consistently bring AI tools into their actual work — not just in low-stakes experiments, but in real

projects where the quality of the output matters. This is how you develop the nuanced judgment that distinguishes effective AI use from superficial familiarity.

The third component is *peer learning*. AI fluency is developing unevenly across professional communities, which means that the colleagues, networks, and communities you engage with can be a significant source of practical knowledge. Seeking out peers who are using AI effectively in your domain — and sharing your own discoveries and failures — accelerates learning in ways that solo study cannot replicate.

Example: A marketing director at a consumer goods company establishes what she calls her “AI Tuesday” habit: every Tuesday morning, she spends thirty minutes reading her curated AI news digest, fifteen minutes experimenting with a new AI tool or feature relevant to her work, and ten minutes sharing one finding with her team in their internal messaging channel. Over six months, this habit — totaling less than an hour per week — has made her one of the most AI-fluent leaders in her organization. More importantly, it has created a culture of ongoing AI learning in her team, with team members regularly sharing discoveries and building on each other’s experiments.

The lesson here is not that you need to dedicate enormous time to AI learning. It is that small, consistent investments in AI literacy compound over time in ways that irregular, intensive bursts of learning do not. The professional who spends one hour per week engaging with AI developments over a year will be far better positioned than the professional who attends a two-day AI conference and then returns to business as usual.

Best Practice

When building your AI news diet, prioritize sources that explain the *implications* of developments for your professional domain over sources that focus primarily on technical architecture. For most professionals, understanding what a new AI capability means for their work is more valuable than understanding how it works under the hood. Recommended starting points include newsletters from reputable technology journalists, industry-specific AI adoption reports from major consulting firms, and curated communities organized around your professional field rather than AI as a general topic.

1.1.6 Setting Your Baseline: An Honest Assessment

Before you can navigate the AI landscape effectively, you need an accurate picture of where you currently stand. This is harder than it sounds. Most professionals either underestimate their existing AI engagement — failing to recognize that many tools they already use are AI-powered — or overestimate their depth of understanding, confusing familiarity with fluency.

A useful framework for self-assessment distinguishes between three levels of AI engagement. *AI-adjacent* describes tools and workflows that are influenced by AI but where you have no direct interaction with AI outputs — your company uses an AI-driven logistics system, for example, and you work with its outputs without thinking about them as AI. *AI-assisted* describes tools where AI enhances or accelerates your work but where you remain the primary author of outputs — a grammar and style checker, a search engine with AI-powered results ranking, or a scheduling tool that suggests meeting times. *AI-powered* describes tools where AI is the primary generator of substantive outputs that you then review, edit, or act upon — a generative writing assistant, an AI research tool, or an agentic system managing a workflow on your behalf.

Most professionals in 2026 are operating across all three levels simultaneously, often without a clear map of which tools fall into which category. Building that map — the personal AI audit described in the exercises at the end of this chapter — is a foundational step. It reveals the true extent of your AI exposure, identifies gaps where AI tools could add value in your current workflow, and surfaces dependencies on AI systems that you may not have consciously recognized.

The goal of this chapter has been to give you the conceptual vocabulary and the contextual understanding to engage with the rest of this book — and with the AI-driven professional landscape — from a position of informed clarity rather than reactive confusion. The AI transition is not something happening to professionals; it is something professionals are actively shaping, one decision, one workflow, and one skill at a time. The chapters ahead will give you the tools to shape it well.

Key Takeaways

- AI has transitioned from experimental technology to an essential workplace infrastructure layer affecting every industry and role. The shift is no longer on the horizon — it is the current reality, and professionals who treat it as a future concern are already operating at

a disadvantage.

- Understanding the difference between narrow AI, generative AI, and agentic AI systems is critical for making informed professional decisions. Each category has distinct capabilities, limitations, and required skill sets, and conflating them leads to both missed opportunities and avoidable errors.
- The pace of AI adoption has created a skills gap that proactive professionals can turn into a competitive advantage. Deep, applied AI fluency remains relatively rare even as surface-level AI use becomes common — which means that genuine competence commands disproportionate recognition and opportunity.
- Staying current with AI developments requires building systematic habits for continuous learning rather than one-time training events. Small, consistent investments in AI literacy — curated reading, active experimentation, and peer learning — compound into significant capability over time.

Exercises

1. **Map your industry's AI adoption curve.** Research three recent case studies of AI adoption in your professional field. For each case study, identify which of the three AI categories (narrow, generative, or agentic) is most prominently featured, what specific business problem the AI is addressing, and what skills the professionals involved needed to work effectively with the system. Use your findings to write a one-paragraph assessment of where your industry currently sits on the AI adoption curve and which category of AI is likely to have the greatest impact on your role over the next two years.
2. **Conduct a personal AI audit.** Create a complete inventory of every digital tool you use in your professional life — including tools you use so habitually that you rarely think about them. For each tool, research whether it uses AI and, if so, categorize your interaction with it as AI-adjacent, AI-assisted, or AI-powered using the definitions provided in this chapter. Reflect on the results: Which tools are you most dependent on? Where are you using AI most deeply? Where might you be missing opportunities to use AI more effectively?
3. **Build your AI news habit.** Identify and subscribe to three reputable AI news sources that are relevant to your professional domain. Spend one week reading them consistently, and

keep a brief daily log of which stories you found most relevant to your work and why. At the end of the week, review your log and use it to refine your source selection — keeping the sources that generated the most relevant signal and replacing those that did not. Commit to maintaining this habit for at least one month before evaluating whether your choices need further adjustment.

PREVIEW

You've reached the end of the pre-view.

The complete edition contains all chapters, including:

- Full chapter content with detailed examples and exercises
- Glossary of key terms
- Appendices and reference material
- A comprehensive index

Get the full book at alkademy.com