

Further Reading

Curated resources organized by pillar. Focused on practical, generalist-accessible content rather than technical papers. Every link has been opened and reviewed — no filler.

Recommended Courses

Free courses that complement this playbook. We've vetted each one for quality, accessibility, and generalist relevance.

- [**AI Fluency: Framework & Foundations \(Anthropic\)**](#) — The academic foundations behind AI fluency. Built on the 4D framework (Delegation, Description, Discernment, Diligence) by Prof. Feller and Prof. Dakan. Free, with certificate. *Best for: understanding the "why" behind AI fluency. This playbook provides the practice; this course provides the conceptual framework.*
- [**Elements of AI \(University of Helsinki\)**](#) — The most widely taken AI literacy course in the world (2M+ learners). No code, no math, no vendor bias. Covers what AI is, what it can and can't do, and how it affects work and society. Created by the University of Helsinki and MinnaLearn. Free, available in multiple languages. *Best for: complete beginners who want a solid, platform-neutral foundation before diving into exercises.*
- [**AI for Everyone 2.0 \(DeepLearning.AI / Coursera\)**](#) — Andrew Ng's updated 2026 version, reframed around decision-making fluency for non-technical professionals. Covers what AI can and can't do, how to evaluate AI projects, and how to think about AI strategy and ethics. Free to audit on Coursera. *Best for: managers, team leads, and generalists who need to make decisions about AI in their organizations.*
- [**Google AI Essentials**](#) — Practical, hands-on introduction with skill badges. Covers prompting, AI tools in the workplace, and responsible use. Leans Google/Gemini for examples, but the concepts transfer. Free. *Best for: people who want to start using AI tools immediately and prefer learning by doing over theory.*
- [**Anthropic's Interactive Prompt Engineering Tutorial**](#) — Hands-on tutorial covering prompt structure, examples, chain-of-thought reasoning, and avoiding hallucinations. Available on GitHub or as a [Google Sheets version](#). Free. *Best for: anyone who wants to immediately improve the quality of their AI interactions.*

By Pillar

Insight Synthesis

- [Anthropic's Prompt Engineering Interactive Tutorial](#) — Hands-on tutorial covering how to structure questions, use examples, and extract better information from AI. The question decomposition techniques directly improve research quality.
- [Prompting Best Practices \(Anthropic Docs\)](#) — The official reference for getting high-quality output from Claude, including techniques for extracting structured information from unstructured sources.
- [Thinking in Systems](#) by **Donella Meadows** — Not about AI, but the best primer on synthesis thinking. Helps you see patterns, feedback loops, and leverage points in any information set.

Workflow Automation

- [n8n Advanced AI Documentation](#) — Reference for building AI-powered automation workflows in n8n. Covers AI agents, LangChain integrations, and connecting AI to your existing tools.
- [n8n AI Workflow Tutorial](#) — Step-by-step beginner tutorial for building your first AI-powered workflow. No code required.
- [The Checklist Manifesto](#) by **Atul Gawande** — Not about AI, but essential reading on why systematic processes beat ad-hoc approaches. Directly applicable to building reliable AI workflows.

Cross-Domain Reframing

- [How Breakthroughs Happen](#) by **Andrew Hargadon** — Research on how innovation comes from recombining ideas across domains, not from inventing from scratch.
- [Range](#) by **David Epstein** — The case for generalists: why breadth of experience produces better problem-solving and more creative AI usage.
- [Prompt Engineering Overview \(Anthropic\)](#) — The official Anthropic prompt engineering guide includes techniques for adapting prompts across different use cases — a practical example of cross-domain thinking.

Agent Collaboration

- [AI Fluency: Framework & Foundations \(Anthropic\)](#) — Free course built on the 4D framework (Delegation, Description, Discernment, Diligence) by Prof. Joseph Feller and Prof. Rick Dakan. The Delegation dimension maps directly to this pillar. CC BY-NC-SA 4.0.

- [Model Context Protocol \(MCP\) Documentation](#) — The open standard for connecting AI to your tools and data sources. If you've completed the advanced agent exercises and want to implement what you designed, start here.
- [The Manager's Path](#) by Camille Fournier — Not about AI, but excellent on delegation, handoffs, and role definition — the same skills you need to orchestrate AI agents effectively.

Ethical Prompting & Judgment

- [On Bullshit](#) by Harry Frankfurt — Short philosophical essay on the difference between lying (knowing the truth and hiding it) and bullshitting (not caring whether it's true). Directly relevant to how AI generates confident but unverified content.
- [Weapons of Math Destruction](#) by Cathy O'Neil — How algorithmic systems can cause harm at scale when deployed without adequate verification. Provides the "why" behind the governance frameworks in the advanced exercises.
- [Anthropic's Claude Model Card](#) — Understanding what the models you're using can and can't do is the foundation of ethical use. The official model documentation is the best source.

General AI Fluency

- [Co-Intelligence](#) by Ethan Mollick — The most practical book on working alongside AI for generalists. Covers mental models, use cases, and limitations without requiring a technical background.
- [XueCodex](#) — A comprehensive technical knowledge base covering AI fundamentals, machine learning, deep learning, and transformers. For readers who want to go deeper into how AI actually works under the hood.
- [The AI Skills Quiz](#) — The quiz that powers this playbook. Discover your pillar scores and learning style archetype.
- [Generalist World Community](#) — Join the conversation with other generalists building AI fluency.

Revision #1

Created 2026-03-16 15:48:27 UTC by Admin

Updated 2026-03-16 15:48:27 UTC by Admin