

Resources

Helpful references to support your AI fluency journey.

- [\[Glossary\]\(/resources/glossary/\)](/resources/glossary/) — Key terms and definitions used throughout the playbook
- [\[Tools and Platforms\]\(/resources/tools-and-platforms/\)](/resources/tools-and-platforms/) — AI tools, agent frameworks, and automation platforms
- [\[Further Reading\]\(/resources/further-reading/\)](/resources/further-reading/) — Curated books, articles, and guides organized by pillar

- [Further Reading](#)
- [Glossary](#)
- [Tools and Platforms](#)

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.

Glossary

Key terms used throughout the playbook, in alphabetical order.

Term	Definition
Agent	An AI system given a specific role, scope, and set of instructions to perform a defined task. In this playbook, "agent" can mean a simple role-prompted AI chat or a more sophisticated autonomous AI system.
Archetype	A learning style profile that describes how you naturally approach new skills. The four archetypes in this playbook (Tinkerer , Planner , Strategist , Social Learner) are derived from AI Skills Quiz results.
Cross-Domain Reframing	The skill of borrowing AI techniques, frameworks, or prompting patterns from one field and adapting them to work in another. One of the five pillars of AI fluency .
Ethical Prompting	The practice of using AI responsibly: verifying outputs, checking for errors, considering audience impact, and maintaining appropriate human judgment. One of the five pillars of AI fluency .
Fluency	In this playbook, fluency means practical proficiency with AI — the ability to use AI tools effectively, critically, and ethically across different contexts. Not just knowing what AI can do, but being skilled at making it work well.
Handoff	The point where output from one AI session or agent becomes input for the next. Managing handoffs — what information to pass, what to leave out, and how to structure the transfer — is a core agent collaboration skill.
Insight Synthesis	The skill of extracting structured, actionable meaning from AI-generated content — distinguishing signal from noise, triangulating across multiple outputs, and forming your own judgment. One of the five pillars of AI fluency .
Multi-Agent System	A setup where multiple AI agents (each with different roles and contexts) work on parts of the same problem, with outputs coordinated by a human or automated orchestrator.
Pathway	A guided route through the playbook's exercises, curated for a specific quiz result profile. Pathways recommend which exercises to do, in what order, based on your strengths and gaps.

Term	Definition
Pillar	One of five core skill areas that make up AI fluency: Insight Synthesis , Workflow Automation , Cross-Domain Reframing , Agent Collaboration , and Ethical Prompting & Judgment .
Prompt Chain	A sequence of AI prompts where each step's output feeds into the next step's input. The simplest form of an AI workflow.
Quality Gate	A checkpoint in an AI workflow where output is evaluated before proceeding to the next step. Prevents errors from propagating through a pipeline.
Tokenization	The process by which AI models break text into smaller chunks called <i>tokens</i> before processing. Tokens can be whole words, parts of words, or punctuation. For example, "tokenization" might become "token" + "ization." Token counts determine input/output limits in AI tools. See How AI Actually Works for more context.

| Workflow Automation | The skill of turning repeatable tasks into systematic AI processes — from reusable prompt templates to multi-step pipelines with quality gates. One of the five [pillars of AI fluency](#). |

Tools and Platforms

All exercises in this playbook work with any AI chat tool. You don't need a specific platform — use whatever you're comfortable with. This page lists options by category for reference.

AI Chat Tools

These are the tools you'll use for most exercises. Any one of them is sufficient.

Tool	Best For	Notes
ChatGPT (OpenAI)	General-purpose exercises, wide model selection	Free tier available. GPT-4o recommended for more complex exercises.
Claude (Anthropic)	Long-form analysis, nuanced reasoning	Strong at synthesis and following complex instructions.
Gemini (Google)	Integration with Google Workspace, multimodal tasks	Good for exercises involving documents or data in Google ecosystem.
Copilot (Microsoft)	Integration with Microsoft 365	Useful if your workflow lives in Outlook, Word, and Teams.

Agent Frameworks

Referenced in the advanced [Agent Collaboration](#) exercises. **Not required for any exercise** — the playbook teaches agent thinking through manual prompting first.

Framework	What It Does	When to Explore
CrewAI	Python framework for orchestrating multiple AI agents with defined roles	After completing AC-Advanced-01 — when you want to automate the handoffs you did manually
AutoGen (Microsoft)	Framework for building multi-agent conversations	When you want agents that can talk to each other without manual copy-pasting
LangGraph (LangChain)	Framework for building stateful agent workflows with graph-based logic	When you need complex conditional logic in your agent pipelines

Automation Tools

Relevant to the [Workflow Automation](#) pillar, especially at the advanced level.

Tool	What It Does	When to Explore
n8n	Open-source workflow automation with AI integration nodes	After completing WA-Advanced-01 — when you want to automate your prompt chains
Make (formerly Integromat)	Visual workflow builder with AI steps	Good for non-technical users who want to automate without code
Zapier	Simple automation connectors between apps, with AI actions	Best for straightforward automations: trigger → AI step → action

Productivity Integrations

Tools that embed AI into existing workflows.

Tool	What It Does
Notion AI	AI writing and analysis built into Notion workspace
Obsidian + Smart Connections	AI-powered linking and search within an Obsidian vault
Raycast AI	Quick AI access from any application on macOS
Google Workspace AI	AI features embedded in Docs, Sheets, Gmail