

Five Pillars

AI fluency isn't one skill — it's a combination of five. We call them ****pillars**** because each one supports the others. You don't need to master all five at once, but understanding them helps you see where you're strong and where you can grow.

These pillars emerged from research and community data gathered through the [AI Skills Quiz](https://aiskillsquiz.com). Each one represents a distinct way of working with AI.

The Pillars at a Glance

Pillar	What It Means	Community Avg
-----	-----	:-----:
[Insight Synthesis](/pillars/insight-synthesis/)	Extracting meaning and patterns from AI output	64%
[Workflow Automation](/pillars/workflow-automation/)	Building repeatable AI-assisted processes	65%
[Cross-Domain Reframing](/pillars/cross-domain-reframing/)	Transferring AI techniques across different fields	67%
[Ethical Prompting](/pillars/ethical-prompting/)	Responsible AI use, verification, and judgment	75%
[Agent Collaboration]		

- [Agent Collaboration](#)
- [Cross-Domain Reframing](#)
- [Ethical Prompting & Judgment](#)
- [Insight Synthesis](#)
- [Workflow Automation](#)

Agent Collaboration

Working effectively with AI agents and multi-agent systems. This pillar covers giving AI defined roles, coordinating between AI sessions, and designing workflows where AI operates with increasing autonomy.

Community average score: 51% — lowest of all five pillars by a 14-point margin. This is the community's biggest growth opportunity, and the area where deliberate practice makes the most visible difference.

Why this pillar matters

Most people interact with AI as a search engine with better grammar — ask a question, get an answer, move on. Agent Collaboration is about treating AI as a team member with a defined role, specific expertise, and clear boundaries.

The shift matters because AI is becoming more capable of multi-step, autonomous work. Tools like Claude Projects, custom GPTs with tools, and AI coding agents already operate at the "agent" level — planning steps, using tools, and making decisions with limited human oversight. Understanding how to collaborate with these systems (not just use them) is the difference between having an assistant and having a team.

The community's 51% average isn't surprising. Most people haven't needed agent-level thinking yet. But as AI tools get more capable, this pillar becomes the differentiator between passive users and active shapers. For a deeper look at the agent spectrum, see [Agents vs. Assistants](#).

What this looks like at each level

Basic — Role-based prompting

You're learning to give AI a defined role and perspective, rather than asking it as a generic assistant. The core skill: framing AI as a specific expert and comparing perspectives to make better decisions.

What it feels like: You run your first "AI team meeting" — asking AI to respond as two different experts on the same problem. You notice that the dual-role output is more nuanced than a single generic answer. You start using role-based prompts by default.

Intermediate — Managing separate agents

You've moved from simulating multiple roles in one chat to actually running separate AI sessions with isolated contexts. Each session has a specific role, limited scope, and doesn't see the other's work. You act as the orchestrator — the person who manages the handoff and synthesizes the results.

What it feels like: You run two AI chats in parallel (strategist and executor), manually pass information between them, and produce a plan that neither could have generated alone. You notice where context gets lost in handoffs and design better transfer summaries.

Advanced — Agent architecture

You're designing complete multi-agent workflows with defined roles, inputs, outputs, handoff triggers, and feedback loops. You think about agent architecture before writing prompts — deciding how to decompose work, what each agent needs to know, and where human checkpoints belong.

What it feels like: You design a 4-agent pipeline (researcher, drafter, critic, editor), run it end-to-end on a real project, and produce output that's measurably better than a single "do everything" prompt. You can explain *why* you split the work the way you did.

Common mistakes

- **"I need agents to be AI-fluent."** No. Most of the value in AI fluency comes from being excellent at the assistant level — writing great prompts, giving AI useful roles, structuring your requests clearly. Agent workflows build on these fundamentals.
- **Over-automating early.** Connecting AI to every tool in your stack sounds powerful, but every connection is a potential failure point. Master manual agent orchestration first, then automate the handoffs.
- **Treating more agents as better.** One well-configured AI with good context will outperform a poorly designed multi-agent system. The value isn't in the number of agents — it's in how well you define their roles and handoffs.

How this connects to other pillars

- [Workflow Automation](#) — agent workflows are the most sophisticated form of AI automation
- [Ethical Prompting](#) — more autonomous AI requires clearer accountability and verification

- [Insight Synthesis](#) — multi-agent outputs need synthesis — comparing, reconciling, and judging across perspectives

Exercises

Level	Exercise	Time	What you'll build
Basic	Your First AI Team Meeting	15 min	A dual-perspective analysis of a real decision
Intermediate	The Handoff Protocol	25 min	A coordinated plan from two isolated AI sessions
Advanced	Design Your Agent Workflow	40 min	A complete multi-agent pipeline with defined handoffs

Cross-Domain Reframing

Applying AI thinking and techniques across different contexts, industries, and disciplines. This pillar is about transferring what works in one domain to solve problems in another — the generalist superpower.

Community average score: 67% — highest of the middle cluster. Users have good instincts here but often lack deliberate practice. Most cross-domain transfer happens by accident; this pillar makes it intentional.

Why this pillar matters

Most people prompt AI using patterns from their own field. Marketers write marketing prompts. Engineers write engineering prompts. Each field develops its own AI patterns — and rarely looks at what other fields have figured out.

But the most powerful AI techniques are often domain-agnostic. A journalist's approach to cross-referencing claims works brilliantly for competitive analysis. An engineer's systematic testing methodology applies perfectly to evaluating AI output quality. A therapist's reframing techniques make excellent prompts for stakeholder communication.

Generalists have a structural advantage here. You work across departments, projects, and contexts. You see how the marketing team's AI challenge is structurally the same as the engineering team's, even though it looks completely different on the surface. This pillar turns that advantage into a deliberate practice.

The connection between cross-domain reframing and AI fluency is this: AI itself is a cross-domain tool. The same model writes code, analyzes poetry, and drafts business strategy. Learning to transfer techniques across contexts mirrors how AI itself works — and makes you dramatically better at using it.

What this looks like at each level

Basic — Borrowing a technique

You're learning to look outside your own field for AI inspiration. The core skill: taking a specific AI technique from an unfamiliar domain and adapting it for your own work.

What it feels like: You discover that data scientists use a particular prompt structure for analysis, adapt it for your project management work, and get a result that's noticeably different from your usual approach. The "stolen technique" reveals that your prompt habits had been constrained by your field's conventions.

Intermediate — Transplanting a framework

You've moved from borrowing a single technique to systematically transplanting an entire problem-solving framework. You map each step of the foreign framework to your context, noting where the mapping is direct, where it needs modification, and where it breaks down entirely.

What it feels like: You take a decision-making framework from military strategy (or medicine, or game design) and apply it to a challenge in your work. The parts where the mapping breaks down teach you more about your problem than the parts where it works smoothly.

Advanced — Building a transfer library

You're systematically collecting, testing, and documenting transferable AI techniques from multiple fields. You build a personal prompt library with tested adaptations, transfer notes, and usage guidance — a resource that compounds over time and becomes shareable.

What it feels like: You have a documented library of 5+ prompt patterns borrowed from other fields. You can articulate *why* a technique transfers, not just that it does. Colleagues start asking you for non-obvious approaches to their AI challenges.

Common mistakes

- **Surface-level borrowing.** Copying a prompt template from another field without understanding the underlying principle produces brittle results. The value is in understanding *why* the technique works.
- **Sticking to adjacent fields.** Marketing people borrow from sales, engineers from product. The most valuable transfers come from genuinely distant domains — the unfamiliarity forces deeper structural thinking.
- **Treating it as a one-time trick.** Cross-domain reframing is a practice, not a hack. The most fluent practitioners make it a habit: every month, explore a new field's AI techniques and test one adaptation.

How this connects to other pillars

- [Insight Synthesis](#) — synthesis becomes more powerful when you can apply insights across different contexts
- [Workflow Automation](#) — the best workflow patterns are often borrowed from other fields
- [Prompt Engineering Basics](#) — cross-domain techniques are fundamentally prompt patterns applied in new contexts

Exercises

Level	Exercise	Time	What you'll build
Basic	The Stolen Technique	15 min	An adapted AI prompt from another field
Intermediate	The Framework Transplant	25 min	A full problem-solving framework adapted for your work
Advanced	The Cross-Domain Prompt Library	40 min	A documented library of transferable techniques

Ethical Prompting & Judgment

Responsible AI use, verification, and critical thinking. This pillar covers knowing when to trust AI output, how to verify it, and how to build systems for accountability.

Community average score: 75% — the highest of all five pillars. Most people *know* they should verify AI output, but there's a consistent gap between "I know I should check" and "I actually have a system for checking." This pillar exists to close that gap.

Why this pillar matters

AI produces confident-sounding text that is sometimes completely wrong. It doesn't flag its own uncertainty, it doesn't distinguish between well-supported facts and plausible guesses, and it will never tell you "I'm not qualified to answer this." That responsibility falls on you.

For generalists, this is especially critical. You work across domains where you're not always the subject-matter expert. When AI generates output about a topic you know well, you can spot errors. When it generates output about a topic you're learning, those same errors become invisible — unless you have a verification process.

The community's 75% average score is encouraging but misleading. People score well on awareness questions ("AI can hallucinate" — yes, most people know this) but lower on practice questions ("I have a systematic way to verify AI output before sharing it"). Confidence without rigor is the most dangerous pattern in AI use.

What this looks like at each level

Basic — Building the verification instinct

You're learning to not accept AI output at face value. The core skill is simple: before using AI output for anything that matters, check it. You're developing the habit of asking "how do I know this is accurate?" after every AI interaction.

What it feels like: You catch your first AI hallucination. You ask the AI to fact-check itself and discover it readily admits to uncertainty when prompted. You build a 3-question verification prompt you start using regularly.

Intermediate — Systematic verification

You've moved from "I should check this" to "I have a checklist for checking this." You've built a verification process tuned to your work — covering factual accuracy, reasoning quality, completeness, and domain-specific concerns. You can evaluate AI output on unfamiliar topics using your process, not just your domain knowledge.

What it feels like: You have a saved verification checklist you actually reach for. You notice when AI reasoning has hidden assumptions. You can explain to a colleague *why* you trust one AI output and not another.

Advanced — Governance and accountability

You're designing systems for teams, not just yourself. You can map AI risk levels across a workflow, create tiered verification processes for different stakes, and write practical guidelines that people actually follow. You think about transparency, attribution, and escalation — not as compliance requirements but as trust infrastructure.

What it feels like: You've written an AI usage guideline that your team references. You've red-teamed your own framework and found the edge cases. You can explain AI-assisted decisions to stakeholders who didn't see the process.

Common mistakes

- **"I always double-check."** Intention isn't process. Without a consistent method, you check when you remember and skip when you're busy. That's when errors get through.
- **Over-verifying low-stakes output.** Not everything needs the same scrutiny. Brainstorming ideas don't need fact-checking; a client-facing report does. Calibrating effort to risk is itself a judgment skill.
- **Trusting AI's self-assessment.** When you ask AI "are you sure?" it will often say yes — or hedge unconvincingly. The Fact-Check Habit exercise teaches you to use structured self-audit prompts that produce genuinely useful uncertainty signals.

How this connects to other pillars

Ethical Prompting is foundational — without it, speed and automation become liabilities. It directly supports:

- [Workflow Automation](#) — every automated AI step needs a quality gate
- [Insight Synthesis](#) — synthesizing AI output requires knowing which parts to trust
- [Agent Collaboration](#) — more autonomous AI requires clearer accountability frameworks

Exercises

Level	Exercise	Time	What you'll build
Basic	The Fact-Check Habit	15 min	A reusable verification prompt
Intermediate	The Verification Checklist	25 min	A systematic verification process
Advanced	The AI Governance Playbook	40 min	A team-level AI usage framework

Insight Synthesis

Extracting meaning, patterns, and actionable insights from AI-generated output. This pillar is about going beyond accepting AI answers at face value — learning to synthesize, compare, and build on what AI produces.

Community average score: 64% — most users are past basic and approaching intermediate. This is the pillar where generalists tend to have natural strengths, because synthesis is fundamentally what generalists do.

Why this pillar matters

AI is excellent at generating volume — 20 ideas in seconds, a 3-page analysis in a minute, a comparison table from multiple data points instantly. But volume isn't insight. The gap between "AI generated a lot of output" and "I extracted something genuinely useful from it" is where this pillar lives.

Most people interact with AI in a single round: ask a question, get an answer, use it or discard it. That's like reading the first page of a research report and calling it analysis. The real value comes from pushing past the first answer — ranking, comparing, contradicting, and finding the pattern underneath.

For generalists especially, synthesis is a superpower. You work across departments, projects, and domains. You're already trained to connect dots that specialists miss. AI amplifies this — giving you more dots, faster. But only if you know how to work with volume instead of being overwhelmed by it.

What this looks like at each level

Basic — Extracting signal from noise

You're learning to take a wall of AI-generated text and pull out what actually matters. The core skill: don't accept the AI's first organization. Impose your own structure — rank by actionability, find the surprising insight, identify contradictions.

What it feels like: You generate a brainstorm of 20 ideas and distill it to the 3 that matter. You notice that the AI organized output by what's common, not by what's useful, and you reorganize it. You start asking "what's missing?" as a default follow-up.

Intermediate — Triangulating across perspectives

You've moved from synthesizing a single AI output to triangulating across multiple AI sessions. You run the same question through different lenses (optimist, skeptic, analyst) and synthesize the results yourself — not by asking AI to do it for you.

What it feels like: You produce a brief that neither any single AI session nor your initial instinct could have generated alone. You notice your own biases in which AI perspective you gravitate toward. Writing the synthesis yourself changes your view of the topic.

Advanced — Research methodology

You're building repeatable research processes with evidence grading, contradiction analysis, and stated confidence levels. You produce decision-ready outputs that distinguish strong evidence from weak, acknowledge genuine uncertainty, and specify what would change your conclusion.

What it feels like: You decompose a complex question into sub-questions, gather evidence systematically, and produce a brief that a decision-maker could act on. You can defend your conclusion *and* articulate what would reverse it.

Common mistakes

- **Accepting AI's ranking.** When you ask AI to rank ideas, it defaults to conventional wisdom. Your domain knowledge and context should override AI's generic prioritization.
- **Asking AI to synthesize for you.** The whole point of synthesis is building *your* judgment. If you ask AI to combine its own outputs, you've outsourced the most valuable step.
- **Stopping at one round.** The first AI output is a starting point. The follow-up questions — "what's missing?", "what contradicts this?", "what's the underlying pattern?" — are where real insight lives.

How this connects to other pillars

- [Cross-Domain Reframing](#) — synthesis becomes more powerful when you apply insights across different contexts
- [Ethical Prompting](#) — knowing which parts of AI output to trust is a synthesis skill
- [Agent Collaboration](#) — multi-agent workflows produce multiple outputs that need synthesis

Exercises

Level	Exercise	Time	What you'll build
Basic	The Signal in the Noise	15 min	A structured insight from a messy brainstorm
Intermediate	The Multi-Source Brief	25 min	A synthesized brief from three AI perspectives
Advanced	The Research Pipeline	40 min	A complete evidence-graded research methodology

Workflow Automation

Building repeatable, AI-assisted processes that save time and reduce manual effort. This pillar focuses on identifying automation opportunities, designing workflows, and integrating AI into existing processes.

Community average score: 65% — solid middle ground. Most users are past basic and approaching intermediate. The gap is usually between "I use AI when I think of it" and "I've designed how AI fits into my recurring work."

Why this pillar matters

Most people use AI in one-off conversations that vanish. They ask a question, get an answer, maybe use it, and start from scratch next time. That's valuable but inefficient — like rewriting the same email from scratch every week instead of using a template.

Workflow Automation is about making your AI usage *systematic*. Instead of ad-hoc prompts, you build reusable templates. Instead of single queries, you chain steps into pipelines. Instead of doing everything yourself, you design processes where AI handles the predictable parts and you handle the judgment calls.

For generalists, this is where AI fluency becomes tangible. You're not just "good at prompting" — you've identified the 3 tasks you do weekly that don't need your brain, and you've automated 2 of them. The time you save goes to the work that actually requires you: empathy, judgment, strategy, relationships.

What this looks like at each level

Basic — Reusable templates

You're learning to capture what works. Instead of typing a fresh prompt every time you summarize meeting notes, you have a template with placeholders: paste in the notes, get a consistent summary. The core skill: separating what stays the same from what changes.

What it feels like: You have a saved prompt template for at least one recurring task. You've tested it with different inputs and refined it. You share it with a colleague and it works for them without explanation.

Intermediate — Prompt chains

You've moved from single prompts to multi-step workflows. You decompose a task into stages (research, draft, refine), give each stage a specialized AI role, and pass outputs between them. The core skill: designing information flow between steps.

What it feels like: You produce a deliverable that went through a 3-step AI pipeline. You notice where context gets lost between steps and design handoffs to preserve it. You document the chain so you can reuse it.

Advanced — Production workflows

You're designing end-to-end processes with quality gates, error handling, and documentation. You map which steps are human vs. AI, define verification checkpoints, and create blueprints that others can run. The core skill: building systems, not just using tools.

What it feels like: You've mapped a business process, identified which steps AI can handle, and built production-ready prompt templates for each. You've measured time savings. Someone else can run your workflow without your involvement.

Common mistakes

- **Automating the wrong things.** The best automation candidates are tasks that are repetitive, predictable, and low-stakes. High-judgment, novel, or high-stakes tasks should stay human-driven (with AI support, not AI control).
- **Skipping quality gates.** Chaining AI steps without verification is like building a pipeline without pressure checks. Each step should have a way to catch bad output before it flows downstream.
- **Over-engineering early.** Start with one reusable template that saves you 10 minutes a week. That's more valuable than an elaborate multi-agent system that you never finish building.

How this connects to other pillars

- [Prompt Engineering Basics](#) — every workflow step is a prompt. Better prompts mean better workflows.
- [Ethical Prompting](#) — automated workflows need quality gates and verification at every stage
- [Agent Collaboration](#) — agent workflows are automation taken to the next level, with AI deciding its own next steps

Exercises

Level	Exercise	Time	What you'll build
Basic	The Reusable Prompt	15 min	A prompt template for a recurring task
Intermediate	The Prompt Chain	25 min	A multi-step AI pipeline
Advanced	The Workflow Blueprint	40 min	A complete AI-automated business process