

The AI Fluency Cheat Sheet

Anthropic's 90-minute AI Fluency course, condensed into 10 pages you'll actually use.

*Prompting is one slice of one competency.
About 25% of using AI well. The other 75% is judgment.*

What's inside:

- The big takeaway: AI skill = 4 competencies, not better prompts
- The 3 modes: Automation, Augmentation, Agency
- The 4Ds in full: Delegation, Description, Discernment, Diligence
- The Description -> Discernment loop (best concept in the course)
- Why AI fails smoothly and how to catch it
- Pre-publish Diligence checklist
- The single best question to ask before shipping AI work

By Kyle Balmer

AI with Kyle | aiwithkyle.com

Original course: Anthropic - AI Fluency: Framework & Foundations (free, 2026)

The Big Takeaway

Most people think AI skill = better prompts. Anthropic disagrees. Their own official course says good AI use is 4 competencies. Prompting is one slice of one of them. Roughly a quarter of the picture.

WHAT MOST PEOPLE THINK

AI skill = better prompts

- Buy prompt packs
- Memorise templates
- Copy other people's prompts
- Tweak adjectives endlessly
- Skip the output review

WHAT ANTHROPIC SAYS

AI skill = 4 competencies

- Delegation
- Description (includes prompting)
- Discernment
- Diligence
- Judgment over hacks

The Course In 30 Seconds

- 1. What AI fluency means** - Using AI well across delegation, briefing, judgment, and responsibility.
- 2. 3 ways people use AI** - Automation, Augmentation, Agency.
- 3. The 4 competencies** - The 4Ds. The framework worth stealing.
- 4. 1 loop that matters** - Description -> Discernment, on repeat.

Good AI use is not just about better prompts.

If you remember one slide from the course, remember this one.

The 3 Modes Of AI Use

Three different interaction modes. Most people get stuck in mode 1 and never move. Mode 2 is where 80% of the leverage lives. Mode 3 is the new frontier.

AUTOMATION *"Do this."*

Give AI a repetitive task. It does it. You walk away.

Best for: transcription, formatting, summaries, repetitive boilerplate.

AUGMENTATION *"Work on this with me."*

You and AI iterating together. Your judgment in the loop constantly.

Best for: drafting, brainstorming, problem-solving, anything creative.

AGENCY *"Go do this and come back."*

You delegate a multi-step outcome. AI plans, executes, returns.

Best for: research, agentic tasks. Claude Code, Codex, Hermes, OpenClaw.

Quick Gut Check

When you opened ChatGPT this morning, which mode did you use? Probably Automation. Probably could have used Augmentation. The leverage compounds in the second mode. Iterate. Don't just dump tasks.

The 4Ds

If you remember one slide from this cheat sheet, make it this one.

D Delegation

What should AI do?

Work allocation. Decide what AI touches before you decide how to prompt it.

D Description

How should I brief it?

Prompting lives here. Context, constraints, examples, format, role. Briefing a collaborator.

D Discernment

Can I trust this output?

The missing half. Evaluate output, thinking, and behaviour. AI fails smoothly.

D Diligence

Did I use it responsibly?

Right tool, right data. Transparency. Verification. The least sexy. Maybe most important.

Description asks well. Discernment judges well.

Most people only do half of either. The 4Ds put the missing halves on the page.

Delegation - What Should AI Do?

The real first question is not 'what's a good prompt'. It's 'what part of this job should the AI even touch'. Most bad AI use is a work allocation problem, not a prompt problem.

The Three Buckets

- ME** Your judgment, context, relationships, accountability.
Your taste. Your call. The things AI shouldn't touch.
- SHARED** Collaborate, iterate, refine together.
Brainstorming, drafting, exploring options - your judgment in the loop.
- AI** Scale, speed, pattern recognition, draft output.
Things you'd do mechanically. AI does them faster and at scale.

Delegation In Real Life

First draft of an email	AI
Final opinion on whether to send it	Probably you
Comparison table of three suppliers	AI
Blindly trusting that comparison table	Never
Hot take on a strategy decision	You
Formatting that hot take into a clean memo	AI
Reading a 100-page report	AI summarises
Forming the conclusion from that report	You

Bad AI use is often a work allocation problem.
Decide what AI touches before you decide how to prompt it.

Description - How Should I Brief It?

This is the slot where prompting actually lives. Anthropic frames it as briefing a collaborator, not summoning magic words.

The Three Ps

Product *What do you want made?*

The deliverable. Format, length, structure, end-state.

Process *How should it approach the task?*

Steps, sequence, what to think about first, what to check.

Performance *How should it behave while collaborating?*

Role, tone, constraints, things to avoid.

The Prompting Checklist

Good prompts cover most of: Context, Examples, Constraints, Steps, Think first, Role/tone.

WEAK PROMPT

Write a marketing plan.

STRONGER PROMPT

Create a marketing plan for the launch of a productivity app.
Target audience: remote professionals.
Constraints: £20k budget, 3-month timeline.
Output: bullet list of key initiatives with brief rationale.

Prompting is Description in practice. Good checklist, not magic.

The RISEN framework (Role/Instructions/Steps/End-goal/Narrow) is a clean way to do it.

Discernment - Can I Trust This Output?

The missing half of how most people use AI. Description asks well. Discernment judges well. Most beginners stop at the first answer. They paste it into the doc and wonder why it sounds AI-generated. (Because it is.)

Three Things To Evaluate

Output *(The result)*

Is the actual answer correct, complete, and useful?

Thinking *(The reasoning)*

Is the logic sound, or just confident-sounding rubbish?

Behaviour *(The collaboration style)*

Is it helpful, or just agreeing with everything you say?

AI Fails Smoothly

AI doesn't crash. It doesn't error. It writes confident, polished, subtly-wrong paragraphs and lets you ship them. Anthropic's example: a paragraph on the Industrial Revolution that reads beautifully. The dates are wrong. The reasoning is shallow. Context is missing. Sounded great. Was wrong.

Polished does NOT mean right.

Confidence does NOT mean competence.

If you skip Discernment, you ship rubbish dressed up as analysis.
Never accept the first answer. Ever.

The Loop That Matters

The single best concept in Anthropic's course. Put it on a sticky note.

Description -> Discernment, on repeat

- 1 DESCRIBE**
Clarify the goal, context, and constraints. Brief the AI.
- 2 INSPECT**
Review the output critically. What's working? What's not?
- 3 TIGHTEN**
Refine the brief. Add clarity, adjust direction, remove noise.
- 4 REPEAT**
Send it back. Get better. Keep evolving.

INTERNET VERSION

Prompt in.
Answer out.
Copy/paste. Done.

REAL VERSION

Brief.
Review.
Redirect. Iterate.

Anyone shipping good AI work is doing this loop. Consciously or not.
If you change ONE thing about how you use AI this week, make it this.

Diligence - Did I Use It Responsibly?

Least sexy section of the course. Maybe the most important. The bit that gets people in trouble - lost jobs, lost clients, sued firms. Don't skip this.

The Three Layers

- Creation** Right tool, right data, right context for the task.
- Transparency** Who needs to know AI helped on this?
- Deployment** Did you actually verify before posting?

The Pre-Publish Checklist

- [] Right tool: not pasting sensitive data into a free chatbot
- [] Right data: not using outdated, incomplete, or off-policy data
- [] Right disclosure: where AI helped on something the audience should know about
- [] Right verification: every claim, citation, stat, name, date checked

THE QUESTION THAT SHOULD STOP YOU SHIPPING

"If this is wrong tomorrow, could I explain what AI did and what I checked?"

If the answer is NO, you skipped Diligence. Don't ship it.
This single question is worth the entire download.

What Anthropic Got Right

Stingy with praise for AI labs. Anthropic earned this one. The course is genuinely more mature than most beginner AI content.

Judgment over hacks

Teaches you HOW to think, not WHAT to prompt. Ages well.

Collaboration over magic

Positions AI as a partner, not a replacement. 'You lead. AI amplifies.'

Human expertise still matters

Your experience is the edge. AI multiplies it. Doesn't replace it.

The Bottom Line

Prompting is the small bit. Judgment is the big one.
Build the 4Ds. Run the loop. Skip the prompt packs.

You lead. AI amplifies.

More from AI with Kyle

Daily newsletter: aiwithkyle.com

Live show: YouTube /@aiwithkyle (Mon/Wed/Fri)

TikTok: @iamkylebalmer

Original Anthropic course: anthropic.com/learn (free)

Source: Anthropic - AI Fluency: Framework & Foundations (2026). Free official course. Condensed for AI with Kyle subscribers. Go to the original if you want the long version.