

TRAINING PREVIEW

Using AI to Think Like a Roman Emperor

Marcus Aurelius had personal tutors who were the finest minds in Rome. You have something comparable in your pocket. Here is how to use AI the way an emperor was taught to think.

Using AI to Think Like a Roman Emperor

March 2026 · For All Staff, Team Leads, Managers

EXECUTIVE SUMMARY

Marcus Aurelius had personal tutors who were the finest minds in Rome. You have something comparable in your pocket. Here is how to use AI the way an emperor was taught to think.

CONTENTS

1. AI Is a Tutor, Not a Search Engine	4
2. Before You Start: What Not to Share	5
3. The Fluency Trap	5
4. Day 1 – Separating Fluency From Depth	6
5. Day 2 – Making AIs Argue With Each Other	7
6. Day 3 – Proving This Works	9
7. From Experiments to Habit	10

Using AI to Think Like a Roman Emperor

Training Preview · March 2026 · For All Staff, Team Leads, Managers

Training preview – a sample of how we teach AI literacy

1 in 4

white-collar workers already use generative AI at work without their employer knowing.^[1] Nearly all of them are using it to go faster. **Almost none of them have discovered what it can actually do for them.**

Jobs and Skills Australia, 2025

This article illustrates the style and structure of AI literacy training we tailor for client workforces. Production versions are adapted to role, risk profile, literacy level, and policy environment.

KEY INSIGHTS

1. **AI is not a search engine.** Used correctly, it is the closest thing to a personal tutor that has ever existed at scale. Almost nobody uses it this way.
2. **Every prompt is a data transfer.** Before you practise, you need to know what not to share. The basics are simple and non-negotiable.
3. **Three days of experiments** you can do at home, on safe topics, that reveal how AI really works, build judgment, and prove the learning is permanent.
4. **The payoff is disproportionate.** The small number of people who do this will have a compounding advantage over those who keep using AI as a shortcut.




Exhibit 1 The three-day learning arc: fluency vs depth, adversarial critique, permanent knowledge, with overnight reflection between each day

Source: Guruswami Advisory

AI Is a Tutor, Not a Search Engine

Most people use AI to get quick answers. That is not surprising. Your brain is built to conserve energy. When a tool offers to do the thinking for you, accepting is the default.

But if you use AI only for speed, you miss what it is actually good at: not giving you answers, but making you smarter. The right kind of AI use forces you to ask better questions, challenge your own assumptions, and notice the gaps in what you thought you knew. It is not pleasant. It requires effort your brain would rather avoid. That is exactly why so few people do it, and exactly why it is an advantage for those who push through.

Marcus Aurelius had this. His personal tutors were among the finest philosophers, linguists, and diplomats the ancient world could offer. That personalised, adversarial, Socratic education shaped the thinking behind Meditations ^[2], a private journal people still read nearly two thousand years later. Not because he was born smarter than everyone else. Because he was taught how to think, personally and relentlessly, by the best minds available.

You have access to something comparable. AI contains the distilled knowledge and reasoning patterns of virtually every field of human learning. It can adapt to your level, challenge your thinking, and surface connections you would not have made alone. The difference between Marcus Aurelius and everyone else in Rome was not intelligence. It was access.

Here is the catch: your brain does not learn at the speed of AI. Learning requires struggle, focused effort on something just beyond your current ability. Your brain flags those active circuits for reinforcement, then does the construction work while you sleep. Spaced repetition, returning to material after gaps rather than cramming, is one of the most robust findings in cognitive science. AI gives you an answer in seconds. Your brain needs days to make it yours. The three days below are designed around this.

This guide is about using AI at home, safely, on topics that carry zero professional risk. Think of it as going to the gym for your judgment.

Before You Start: What Not to Share

Every prompt you type leaves your device and is processed by an organisation whose interests are not yours. Over weeks, your prompts build a profile of what you are working on, what problems you are solving, and what decisions you are weighing. Individually harmless. Collectively revealing.

On personal accounts at home, never share live company strategy, client names, financial results, internal incidents, or anything you would not post on LinkedIn. On work systems, follow your corporate AI policy. If there is no policy, treat every public AI tool as if it publishes your input.

Now, an honest moment. Your boss sends you a task at 7PM. You are tired, the response is urgent, and you need to summarise a document or draft a quick reply. You reach for the AI tool. You know the policy. You convince yourself the risk is small.

Please do not. I have investigated hundreds of security incidents over thirty years. The ones that cause the most damage almost never start with malice. They start with a tired person making a reasonable decision under pressure. The data leaves the network. It does not come back. For the full picture on where corporate data goes, see our brief on [shadow AI and data leaving your network](#)^[3].

The Fluency Trap

AI collapses the time it takes to sound knowledgeable. You can go from knowing nothing about a topic to speaking fluently about it in minutes. The vocabulary is right. The structure is logical. The confidence is high. But fluency is not depth. This is the [Fluency Trap](#)^[4], and it affects everyone.

The good news: you can train yourself to spot it. Set aside 30 minutes when your brain is fresh and you are genuinely curious to learn. Not when you are tired, distracted, or rushing through your evening.

You are about to see something few people know, and it will teach you more about how AI and you should interact than any PowerPoint slide can show. These experiments are fun. You will laugh, you will learn, and you may even be horrified at what you discover.

Day 1 – Separating Fluency From Depth

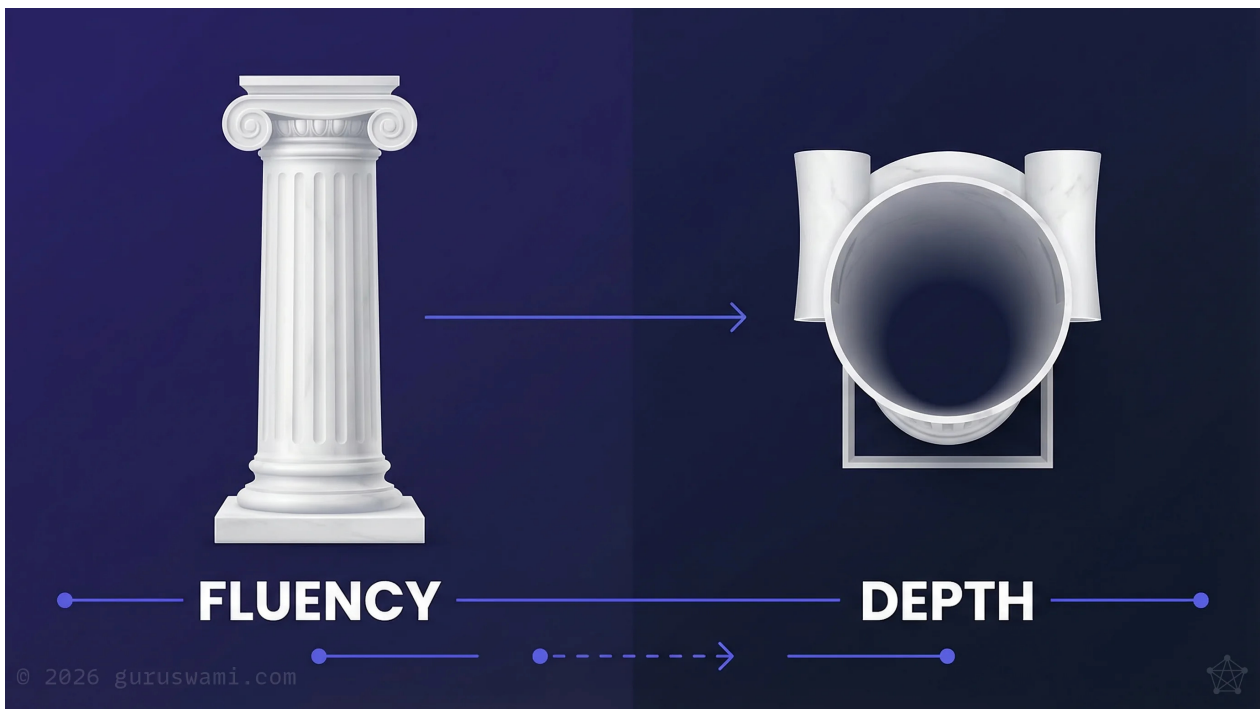


Exhibit 2 The fluency trap: what AI sounds like versus what AI actually knows

Source: Guruswami Advisory

This experiment uses AI to show you exactly how the fluency trap works. Do it on a topic you already know well so you can judge the quality of what comes back.

Step 1. Pick something you genuinely understand. Not work-related. Something personal and harmless: a sport you play, a hobby you have spent years on, an industry you worked in previously, a subject you studied. The key is that you must be able to tell when the AI gets it wrong.

Step 2. Write a short paragraph stating your view on something specific within that topic. A technique, a common misconception, a trade-off that most people get wrong. Be specific.

Step 3. Give your paragraph to the AI and ask:

"Make the strongest case that this view is wrong. Challenge my assumptions, identify what I might be missing, and point out where my reasoning is weakest."

Then separately ask:

"Now make the strongest case that this view is right. Provide the best evidence and examples that support it."

Step 4. Read both responses carefully. Compare them against what you actually know. Ask yourself:

- Where did it reveal ideas I had not considered? (This is AI at its best.)
- Where did it sound confident but get things obviously wrong? (This is the fluency trap in action.)
- Did it invent facts, cite sources that do not exist, or present opinions as evidence?
- Would someone who does not know this topic be able to tell the difference?

That last question is the one that matters. If someone without your expertise would read the AI's response and believe it completely, you have just seen exactly how the fluency trap works in every domain where you are not the expert.

Now stop. You have done 30 minutes of focused work and your brain has plenty to process. Let the results sit for a day. You will find yourself thinking about what the AI got wrong, what it got right, and what you were not as sure about as you thought. That reflection is not idle. It is part of the learning.

Come back tomorrow for Day 2.

Day 2 – Making AIs Argue With Each Other

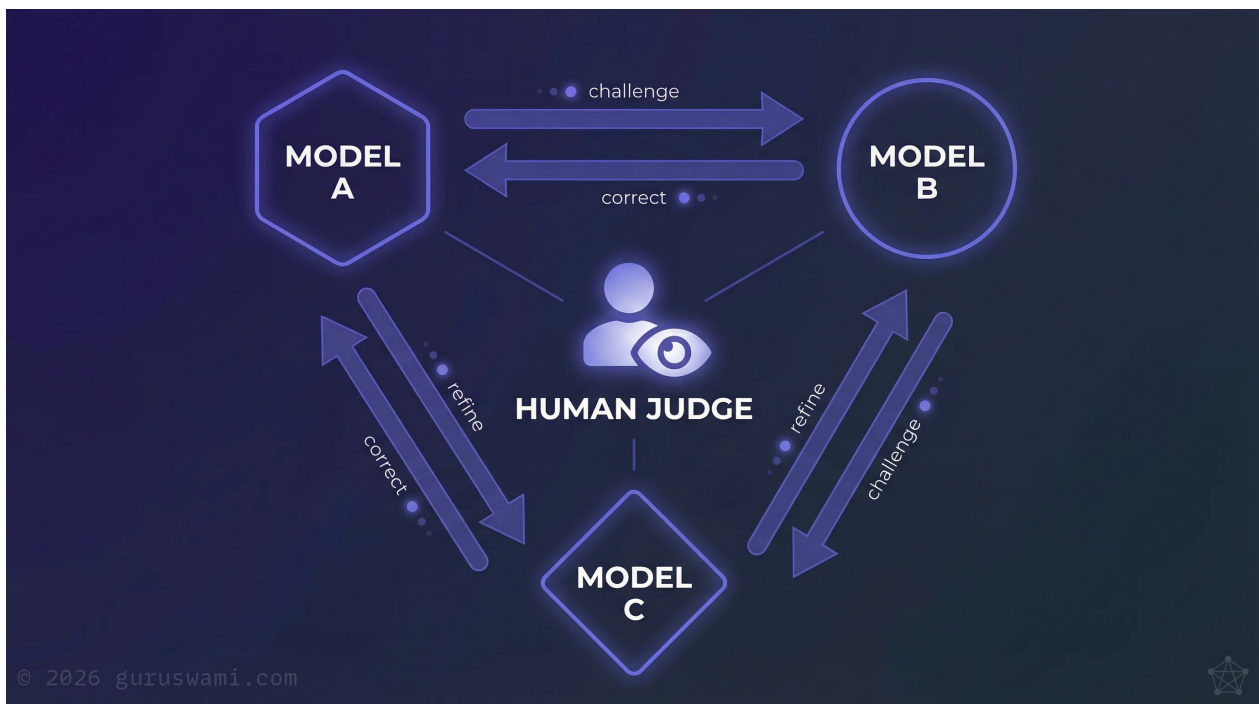


Exhibit 3 Adversarial cross-model critique: three AI models challenging each other with a human judge at the centre

Source: Guruswami Advisory

Different AI tools give different answers to the same question. They have different training data, different biases, and different failure modes. When you make them critique each other's work, something happens that goes beyond what any single tool can produce.

Step 1. Sign up for two or three free AI tools. ChatGPT, Gemini, Perplexity, Claude. Most have free tiers. Use your personal email, not your work account.

Step 2. Use the same safe, familiar topic from Day 1.

Step 3. Ask each tool the same question: your view, plus the request to challenge it (exactly as in Day 1).

Step 4. Now make them talk to each other. Take the critique from one model and paste it into another, along with the prompt below. Then press enter, sit back, and be prepared for the fun to begin.

"Here is a critique of a position I hold, written by a different AI. Evaluate this critique. Where is it strong? Where is it weak? What has it missed? Where has it stated something as fact that is actually wrong?"

Step 5. Take that response back to the first model. Repeat. Iterate until the models reach some form of consensus, or until they diverge dramatically. Both outcomes are valuable.

You are now watching the world's best billion-dollar AIs argue with each other, over your favourite topic.

What you will see:

- The models improve each other's work. Each round of adversarial revision sharpens the analysis.
- They catch each other's hallucinations. You will see admissions like "upon review, the citation I provided does not appear to exist."
- They sometimes converge on things that are still wrong. Consensus between AI models is not truth. It is the intersection of their training data.
- The quality of the final output is dramatically higher than any single tool produced alone.

You, the person with actual expertise, are the only one who can judge what is real. That judgment is exactly what this experiment builds. This is the Mirror Principle^[5] in action: AI reflects human knowledge and bias, not objective truth. Different models are different mirrors. When you make them reflect off each other, you get a richer picture, but the person looking into them remains the only one who can tell what is real.

Stop again. Another 30 minutes of intense focus, another day to let it settle. You will notice things tomorrow that you missed today. That is the point.

Day 3 – Proving This Works

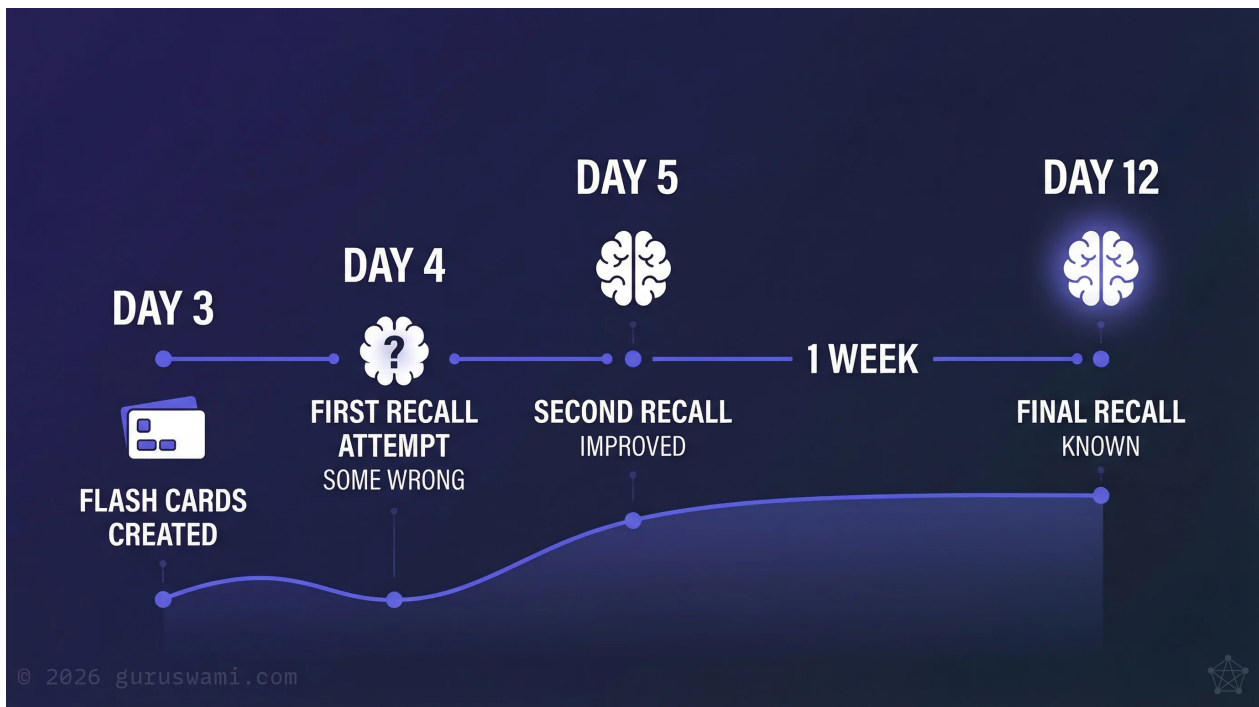


Exhibit 4 Spaced retrieval timeline: flash cards on Day 3, first recall on Day 4, improved recall on Day 5, permanent knowledge by Day 12

Source: Guruswami Advisory

This is the biggest payoff. Days 1 and 2 showed you the limits of AI and the limits of your own expertise. This one proves you can use both to actually learn something permanent.

Step 1. From your adversarial sessions, identify five ideas that were genuinely new insights into your chosen topic. Not things you already knew. Not things the AI got wrong. Five things that survived your scrutiny and expanded what you understood.

Step 2. Ask whichever AI you found most capable to help you learn these topics. Ask it to create flash cards, prompt questions, and short summaries of the major points. Be specific about what you want to retain.

Step 3. Print them out. Physical paper, not a screen. Read them tomorrow, when you are fresh and in a learning mood. For each card, attempt to recall the answer before you flip the page and read the complete answer. Reflect on your mistakes. Acknowledge you got some wrong. Then stop thinking about it entirely.

Step 4. Repeat this exercise the following day, ideally after a good night's rest. Review the results. You did better this time.

Step 5. Test again in one week. The answers were easy. You got pretty much everything correct. But something new happened. These answers felt known, familiar, and they prompted you to think further, to ask even more questions.

That is spaced retrieval, sleep consolidation, and the compounding effect of genuine learning. You did not just read about it. You felt it happen.

The five things you learned are now yours. Not the AI's output. Not a summary you will forget by next week. Actual knowledge, integrated into the way you think about a subject you already cared about. And you built it in under a week, using free tools, on a topic that carries zero professional risk.

Now imagine doing this deliberately, once a month, on topics relevant to your work. Not pasting documents into AI. Not asking it to do your job. Using it to make you better at your job. That is the difference between AI as a shortcut and AI as a tutor.

From Experiments to Habit

Keep a learning log. After each session, write down one thing that surprised you or that you later checked and found wrong. Over weeks, you build a personal record of where AI is reliable and where it is not.

Do this once a month. Pick a new topic, or go deeper on an existing one. Each cycle sharpens your judgment and builds on the last.

KEY TAKEAWAYS


- AI is a tutor, not a shortcut. Used correctly, it forces you to think harder, not less. Almost nobody uses it this way.
- Every prompt is a data transfer. Keep work content off personal AI accounts. If your employer has no AI policy, treat every public tool as if it publishes your input.
- The fluency trap is real. AI sounds right regardless of whether it is right. Day 1 lets you see this firsthand on a topic where you can judge quality.
- Different AI tools give different answers. Making them critique each other produces sharper analysis than any single tool alone.
- Spaced retrieval works. Day 3 proves it on yourself: five new insights, flash cards, three sessions over a week, and the knowledge is yours permanently.
- When your AI use starts involving real work decisions, push your organisation to provide governed tools. Not personal accounts.

If you did these experiments and discovered something new, we would love to hear what you experienced. The discovery of this was profound for us too. Share it with your friends and family. Encourage them to try it. Imagine what it does when it becomes a regular habit.

We deliver structured AI literacy workshops adapted by audience, from board-level briefings to hands-on team sessions. Each is tailored to your organisation's risk profile, policy environment, and current capability. This article is a sample of the approach. For the executive perspective, see our [brief on AI Literacy as an Executive Advantage](#)^[6].

References

1. <https://www.jobsandskills.gov.au/publications/generative-ai-capacity-study-report>
2. <https://en.wikipedia.org/wiki/Meditations>
3. <https://guruswami.com/insights/shadow-ai-data-leaving-your-network/>
4. <https://guruswami.com/insights/ai-literacy-executive-advantage/#the-fluency-trap>
5. <https://guruswami.com/insights/ai-literacy-executive-advantage/#the-mirror-principle>
6. <https://guruswami.com/insights/ai-literacy-executive-advantage/>



About Guruswami Advisory

Independent AI security and strategy advisory for Australian boards, leadership teams, and regulated organisations. No vendor ties. No platform allegiance. Every recommendation tested on our own infrastructure.

Paul Nevin, Principal Advisor. 28 years in cybersecurity and cyber-intelligence. Six years of applied AI research. Every engagement led personally.

Contact

info@guruswami.com

guruswami.com

[linkedin.com/in/paul-nevin](https://www.linkedin.com/in/paul-nevin)

Guruswami™ Pty Ltd | ABN 11 695 354 020 | Canberra, ACT, Australia

This document is provided for informational purposes. It does not constitute legal, financial, or insurance advice. Where findings have regulatory implications, engage qualified legal counsel.