

# SAGE

*Structured AI-Guided Education*

---

## **Your Guide to Using AI Responsibly**

*Five steps that make your work genuinely yours*

**Developed by Dr Mahmoud Elkhodr & Professor Ergun Gide**

*Central Queensland University · [sage-framework.com](https://sage-framework.com)*

## What is SAGE?

SAGE is a five-step process for using generative AI tools in your coursework and assessments. It is not about whether you can use AI — your lecturer sets that rule. It is about how to use it so that the thinking in your work is genuinely yours.

AI tools are fast and fluent. They can produce a convincing-looking answer to almost any question in seconds. That is the problem. Fluent does not mean correct, and correct for one context does not mean appropriate for yours. SAGE gives you a repeatable process to catch what AI gets wrong, add what it misses, and document the reasoning that shows you actually understand your work.

The short version: generate a draft with AI → check it against real sources → fix what is wrong → stress-test your fixes → reflect on what you learned. Each step leaves a record of your thinking. That record is your work.

## Before You Start: Find Your Anchors

Every step in SAGE involves checking AI output against something authoritative. That authoritative source is called an anchor. Your lecturer will usually provide your anchors — a marking rubric, a required reading, a professional standard, or a regulatory framework.

If your lecturer has not specified anchors, you need to find them yourself. This is a real professional skill. In your career, you will always need to know what the authoritative reference is for your field before you can evaluate whether AI output is reliable.

Your discipline	Where to find your anchors
<b>Cybersecurity</b>	NIST Cybersecurity Framework, ISO 27001, Privacy Act 1988, ASD Essential Eight, ACSC guidelines
<b>Health / Nursing</b>	Clinical practice guidelines, NDIS Practice Standards, NSW Health policies, AHPRA codes, drug reference compendia
<b>Business / Management</b>	Relevant legislation (Fair Work Act, Corporations Act), company policies, ASX governance codes, industry codes of conduct
<b>Engineering</b>	Australian Standards (AS series), Safe Work Australia codes, engineering codes of practice, environmental regulations
<b>Law</b>	Relevant statute, case law, Legal Aid guidelines, court rules, professional conduct rules
<b>Education</b>	Australian Curriculum (ACARA), state syllabus documents, UDL guidelines, school/institution policy
<b>Not sure?</b>	Start with your unit reading list. Find the key textbook or foundational paper for your topic. Ask your librarian for the peak professional body in your field.

### How to verify a source is authoritative

Ask: Is this published by a government body, a peak professional association, a regulatory agency, or a peer-reviewed journal? If yes, it can be an anchor. A blog post, a Wikipedia article, or another AI-generated summary cannot.

## Step 1

**1 Generate***Start with your thinking — not the AI's.*

Before you type a single word into an AI tool, spend five to ten minutes writing down what you already know about the task. It does not need to be polished. Bullet points, rough notes, or a badly worded draft all count. The point is that your thinking comes first.

Then build your prior knowledge into your prompt. Give the AI context about the task, the constraints, and the standards that apply. A prompt that includes what you already know will produce far more useful output than a blank question.

**💡 Why this matters**

AI produces generic output in response to generic prompts. When you include your own context — your field, your constraints, your standards — the AI has something to work with. You also immediately see where the AI diverges from what you already knew, which is exactly what you need for Step 2.

Use your lecturer's base prompt if one has been provided. If not, use this structure:

**EXAMPLE PROMPT**

*I am a [year level] [discipline] student. The task is: [describe the task in your own words]. The relevant standard or framework for this context is [anchor]. My initial thinking is: [paste your notes or draft here]. Based on this context, [specific request].*

**💡 What if you have no prior knowledge at all?**

Write down what you understand the task to be asking, which terms you recognise, and where your understanding breaks down. Even a statement of what you do not yet know is a valid starting point — and it is far more useful to include in your prompt than nothing.

## Step 2

**2 Evaluate***Check the AI's work against something real.*


Read the AI output carefully. It will probably look confident and well-structured. That does not mean it is right. AI tools frequently miss context-specific requirements, omit regulatory details, produce plausible-sounding but incorrect facts, and apply generic advice where specialist knowledge is needed.

Compare the AI output against your anchors point by point. For each claim the AI makes, ask:

- Is this accurate according to my anchor?
- Is this complete, or has the AI missed something important?
- Is this appropriate for my specific context, or is it generic advice that does not fit here?

→ Has the AI omitted any regulatory, ethical, cultural, or stakeholder requirements?

Take notes on everything you find. These notes become your decision log in Step 3.

 **Common AI failure patterns to look for**  
 AI frequently: omits enforcement mechanisms; misapplies standards from a different jurisdiction; gives generic best practice rather than context-specific requirements; misses exception handling and edge cases; and ignores cultural, accessibility, or stakeholder-specific factors.

**3** Step 3 **Refine** *Fix it, own it, document every decision.*

Now revise the AI output based on what you found in Step 2. For every significant element, make a deliberate decision and record it in your log. There are only three options:

**ACCEPT** The AI output is accurate and appropriate for your context. You can explain why, citing your anchor.

**MODIFY** The AI output is partially right but needs adjustment. Record exactly what you changed and why, citing the specific source or reasoning.

**REJECT** The AI output is wrong, misleading, or inappropriate for your context. Record what you removed and what the correct position is, with a reference.

Use the table below to record your decisions. Your log is the core evidence that the thinking in your work is yours. Keep it.

AI Output — paste or summarise the claim	Decision	Justification — cite a specific source, standard, or reason (50–75 words)
<i>Example: "AI recommended optional encryption for stored data"</i>	<i>Reject</i>	<i>Encryption must be mandatory for health records under Privacy Act APP 11.1. "Reasonable steps" threshold is not met by optional controls. Source: OAIC guidance on APP 11.</i>

--	--	--

**💡 How many decisions should you log?**

Aim for at least one logged decision per significant paragraph or section of the AI output. A surface-level assessment with only two or three entries does not demonstrate genuine critical engagement. Six to ten decisions for a standard-length task is a reasonable target.

**4 Step 4  
Audit**

*Use AI to stress-test your reasoning — then judge the AI.*

You have made your decisions and documented your reasoning. Now use AI to challenge you.

Write a new prompt that tells the AI to act as a professional expert matched to the anchor you used in Step 2. Give it your decision log and ask it to find weaknesses in your reasoning.

**EXAMPLE PROMPT**

*You are a [professional role — e.g., senior cybersecurity auditor / clinical reviewer / compliance officer]. I have made the following decisions when evaluating AI-generated content against [your anchor]. Review my log and identify: where I may have accepted something I should have questioned; where a justification is weak or missing a citation; and where my reasoning may not hold in this professional context. [Paste your decision log here]*

The AI will flag concerns about your decisions. Your job is to evaluate those flags — not just accept them. Record your adjudication in the table below.

AI flag — what the AI criticised	Your decision	Your reason — why you accept or reject this flag (30–50 words)

**This step is a single bounded pass. You do not keep going back and forth with the AI. One set of flags, one set of adjudications, documented. You are the decision-maker — the AI is the challenger.**

**💡 When to reject an AI flag**

Reject the flag when: the AI is applying a generic principle that does not fit your specific context; the AI does not know about a stakeholder constraint or local requirement you do know; or the AI is repeating a concern you already addressed in your log. Always record why you are rejecting it.

Step 5  
**5 Reflect**

*What did you actually learn?*

Look back over your decision log and your audit adjudications. This is where you turn the experience into learning that will carry forward beyond this task.

Answer the four questions below. Your reflection should be specific — not general observations about AI, but concrete observations about what this particular AI tool did in this particular disciplinary context.

Reflection question	Your answer
<i>What did the AI consistently get wrong in this task?</i>	
<i>Where did your own knowledge prove essential — something the AI could not have known?</i>	
<i>What is the right "anchor" for this kind of task in a professional context after graduation?</i>	
<i>What would you do differently next time?</i>	

The habit SAGE builds is not "check AI with these specific sources." It is "identify the right anchor for any context, then hold AI accountable against it." In your career, the anchors will change with every role and every task. The judgement stays the same.

 **If your lecturer includes a supervised session**

Your decision log and reflection are your preparation. The session is not a test of memory — it is a conversation about your reasoning. If you can walk through your log and explain why you made each decision, you are ready.

#### ATTRIBUTION & LICENCE

Developed by Dr Mahmoud Elkhodr and Professor Ergun Gide, Central Queensland University. Validated across 800+ students. Listed in the TEQSA Generative AI Knowledge Hub. CC BY 4.0 — free to use and adapt.

If your lecturer has adapted this guide: *Adapted from the SAGE Framework (Elkhodr & Gide, 2026). sage-framework.com*