

MODULE 4 · From AI-Proofing to AI-Resilience

Lecture 4.3

Using AI as Tutor, Adversary, or Collaborator



If you do not specify, you get Oracle.

When “AI permitted” is left unspecified, students may reasonably treat AI as an answer generator. Specify the intended posture.

Match posture to AIAS level.

Posture	AIAS levels it fits	AI Action	Evidence	Typical course context
Tutor	Levels 2–3	Asks questions; withholds answers.	Dialogue + reflection.	Introductory or intermediate
Adversary	Levels 3 or 5	Generates flawed outputs.	Student critique.	Intermediate or advanced
Collaborator	Levels 3–4	Student drafts first; AI suggests structure.	Final draft + decisions note.	Advanced or applied

Posture mismatch is the most common design error.

Select the posture based on the cognitive operation students must perform, not convenience.

Tutor, specified.



STUDENT-FACING

You will work with AI as a Tutor. The Tutor's instruction is included below. The Tutor will ask you questions; **it will not give you the answer**. After at least two attempts, the Tutor may give a hint. Your submission is the dialogue (copy-paste) plus a 100-word reflection on what shifted in your thinking.

Tutor boundary: No answer before two documented attempts.

Adversary, specified.



STUDENT-FACING

Use the prompt below to have AI generate three plausible-but-flawed responses to **[task]**. Write a 200-word critique of each response, identifying the flaws and explaining the consequence for **[domain]**. **Your submission is the critique, not the AI output.**

Adversary boundary: The student submits the critique, not the AI output.

Collaborator, specified.



STUDENT-FACING

You will work with AI as a Collaborator. **Write your draft first (no AI).** Then use AI to suggest structural changes. You decide which suggestions to accept. Document your decisions in a 100-word note: "I accepted [X] because __; I rejected [Y] because __." Submission: final draft + decisions note.

Collaborator boundary: The student drafts first and records accepted and rejected suggestions.

iCRAFT structure.

I	Instruction	"Generate three plausible-but-flawed responses to ____."
C	Context	"This is for a sophomore writing course on persuasion."
R	Role	"Act as a peer student."
A	Audience	"For your classmates in this course."
F	Format	"Three short paragraphs, each labeled."
T	Tone	"Plain academic; no hedging."

Use iCRAFT only when students are expected to run an instructor-provided AI prompt.

The posture-preserving boundary.



Every AI-integrated assignment must include one boundary that, if violated, signals the posture has collapsed. **Example:** *"If the AI gives the answer before you have shown two attempts, paste a screenshot and explain what happened. This is documentation, not penalty."*

Boundaries that punish drive students underground. Boundaries that document teach calibration.

This is documentation, not penalty. Boundary = trigger + documentation response
+ restorative follow-up.

Free-tier check.

Before you finalize your AI-integrated assignment, test that it can be completed on the free tier of a major LLM. If it cannot, choose a different posture, simplify the prompt, or provide a documented non-AI alternative. The non-AI alternative must assess the same learning outcome without adding extra work or penalty.



Access risk *Test on a free tier before you ship. A paywalled assignment is an access barrier.*



Elaborate *In one sentence, name the AI posture your redesigned assignment will use and the boundary that keeps it intact.*

The AI-Resilience Checklist.

Phase 1 (Alignment): Does the learning goal reinforce the AIAS level and evidence form?

Phase 2 (Evidence): Is there exactly ONE piece of process evidence? Does it pass the workload threshold?

Phase 3 (Posture): If AI is integrated, is the role specified (Tutor/Adversary/Collaborator) and verified on a free tier?