When to Build In (or Keep Out) AI

A Decision Framework for Assignment Design

Center for Teaching, Learning, & Outreach

The Core Challenge

AI tools can either accelerate learning or short-circuit it — and the difference depends almost entirely on assignment design, not on blanket policies. This framework helps you make a decision assignment-by-assignment, grounded in what we know from learning science.

When AI helps learning

- Offloads rote/automatable subtasks
- Mirrors how work is done professionally
- Frees cognitive load for higher-order thinking
- Teaches discernment and critical evaluation
- Scaffolds without replacing core reasoning

When AI harms learning

- Bypasses the building of mental models in novice learners
- Removes desirable difficulty before mastery
- Replaces the reasoning the course is building
- Produces confident-but-wrong answers students accept
- Undermines certification/assessment of individual competence

Theoretical Foundations

The decision tree on the following pages is grounded in five learning science frameworks. Understanding these helps you reason about novel situations the tree does not cover.

Theory	Implication for AI policy
Transfer-Appropriate Processing	Skills are best retained when practice conditions match retrieval conditions. Foundational skills used without AI in the field should often be learned without AI — at least until fluency is established.
Cognitive Load Theory	Novices need to struggle productively to build mental models. AI removes that struggle before the model exists, producing "hollow fluency" — students who can present but not reason.
Bloom's Taxonomy	AI can handle Remember and Understand tasks with ease. The real learning (Analyze, Evaluate, Create) is where AI should be a collaborator, not a substitute. Design rubrics that weight the top levels heavily.
Situated Learning	Learning is most effective when it mirrors authentic practice. At the advanced stage, restricting AI may produce graduates who are less competitive because they have never learned to work alongside AI as practitioners do.
Critical Evaluation of AI Output	AI confidently produces plausible-but-wrong output. Students without sufficient domain expertise cannot detect these errors — they will learn false things. Build domain knowledge first; introduce AI tools after students can audit their output.

The Decision Tree: Seven Questions

Work through these questions (in order) for any assignment you are designing or reviewing. Each question routes you toward one of eight outcome policies (see next section).

Key question	Theory	Yes →	No →
Q1. Is this a foundational disciplinary skill — one students must perform independently as practitioners?	*Transfer-Appropriate Processing*	**Go to Q2**	**Go to Q3**
Q2. Is the student still in the early/novice phase — building mental models from scratch?	*Cognitive Load Theory*	**Go to Q2a**	**Go to Q2b**
Q2a. Does the assignment require original reasoning, argument, or analysis?	*Bloom's Analyze / Evaluate*	**Restrict (A)**	**Hybrid scaffold (B)**
Q2b. Would AI use mirror how this skill is actually used professionally?	*Situated Learning*	**Allow — collab (C)**	**Go to Q4**
Q3. Is the primary learning goal a process (thinking, iteration, revision) rather than a polished product?	*Metacognition / Formative*	**Go to Q3a**	**Go to Q5**
Q3a. Can you make the thinking process visible and assessable?	*Zone of Proximal Development*	**Hybrid — transparent (D)**	**Restrict — process (E)**
Q4. Could AI produce convincingly wrong answers that students lack the expertise to detect?	*Epistemic Vigilance*	**Restrict — epistemic (F)**	**Allow — productivity (G)**
Q5. Is this a summative/high-stakes assessment that must certify individual competence?	*Assessment Validity*	**Restrict — authentic (H)**	**Go to Q6**
Q6. Would AI free cognitive resources so students focus on higher-order aspects of the task?	*Bloom's Create & Evaluate*	**Allow — higher order (I)**	**Restrict — rote (J)**

Eight Outcome Policies

Each outcome includes the policy name, practical tactics, and typical examples.

Policy	Tactics	Typical examples
A — AI-Resistant (core skill)	**Protect foundational skills from AI substitution** – In-class writing, oral exams, lab practicals, timed problem sets – Exam wrappers: reflect on reasoning process after the test – Partial credit for visible reasoning/show-your-work – Paired verbal defense: explain any written submission	*First-year proof writing, introductory clinical reasoning, foundational coding logic, early legal analysis*
B — Bounded AI use	**Allow AI for surface tasks; protect the cognitive core** – Specify permitted uses (e.g. proofreading, not drafting) – Require an AI use log: what was generated vs. the student's own – Process portfolios with multiple drafts and reflections – Assign the AI editor role, not the AI author role	*Research papers (analysis must be original); lab reports (AI may format results tables, not interpret findings)*
C — AI-Integrated (professional)	**Teach students to collaborate with AI as practitioners do** – Explicitly teach prompt engineering and iterative refinement – Require students to critique AI output: what did it get wrong? – "AI + human" projects with clearly bounded student contribution – Evaluate quality of AI collaboration, not just the final product	*Advanced research synthesis, professional writing (grant proposals, reports), software development, data analysis pipelines*
D — Transparent AI use	**Allow AI, but assess the thinking — not just the product** – Annotation: students annotate AI text explaining edits/decisions – Reflective memo: what I asked, what it gave me, what I changed – Draft comparison: AI draft vs. final with justification for all edits – Process portfolios with timestamped drafts and self-assessment	*Writing assignments with required revision histories; design projects with documented iteration*
E — Restrict (process)	**Redesign to make process assessable, or restrict AI fully**	*Problem-solving assignments where only the answer is submitted; take-home essays with no drafts required*

Policy	Tactics	Typical examples
	– Move to live, in-person formats where process is observable – Replace take-home with in-class versions of the task – Oral exams or presentations: explain and defend thinking live – Thinking-out-loud audio/video recordings during task completion	
F — Restrict (epistemic risk)	**Build domain expertise before introducing AI tools** – "Spot-the-error" exercises using AI-generated content to critique – Fact-verification labs: check AI claims against primary sources – Introduce AI tools only after students demonstrate baseline competency – Assign an AI-skeptic role: challenge every AI claim in group work	*Medical diagnosis, historical analysis, legal reasoning, scientific interpretation — domains where AI confidently produces plausible-but-wrong content*
G — Allow (productivity)	**Students can evaluate AI output and are developing professional judgment** – Require a brief AI use statement noting what was used and verified – Grade quality of AI-human synthesis, not AI use per se – Use AI-generated drafts as starting points for critical revision – Discuss AI output in class as live examples of good/bad responses	*Advanced literature reviews, code refactoring, data cleaning, draft generation the student then substantially revises*
H — Authentic assessment	**Require authentic, individual demonstration of competence** – Oral examinations and viva voce defenses – In-class timed assessments (papers, cases, problems) – Portfolios with required in-person discussion of selected pieces – Idiosyncratic prompts tied to each student's prior work or context	*Capstone defenses, clinical competency assessments, comprehensive exams, final essays requiring integration of personal fieldwork*
I — Strategic AI integration	**Use AI to free cognitive runway for higher-order thinking** – Give students an AI draft; ask them to evaluate, revise, and defend it – Use AI for research aggregation; require students to synthesize/argue – Let AI handle formatting/citations; students focus on argumentation – Design rubrics weighting synthesis and judgment heavily	*Policy briefs (AI compiles background; student argues), design projects (AI generates options; student selects and justifies), advanced seminars*

Policy	Tactics	Typical examples
J — Restrict (productive struggle)	**The rote work is the learning — protect it** – Be explicit with students about why the "tedious" work matters – Design so the rote work reveals misconceptions – Use low-stakes AI-free tasks as diagnostics before high-stakes ones – Acknowledge this is temporary: "once you've internalized this, AI becomes useful"	*Manual data-entry exercises that build schema understanding, handwriting-heavy first drafts that slow thinking productively, step-by-step proofs before using CAS tools*

Design Tips for Any AI Policy

Regardless of where your assignment lands in the tree, these practices improve the integrity and learning value of the work.

If you are restricting AI

- Make the restriction purposeful and communicate the learning rationale to students
- Redesign delivery so restriction is inherent (in-class, oral, timed)
- Build in formative practice before high-stakes AI-free assessments
- Treat it as temporary: plan when you will introduce AI tools as students develop mastery

If you are allowing or integrating AI

- Specify which tasks AI may and may not do — in writing, in the assignment prompt
- Require an AI use statement or process artifact that makes AI contribution visible
- Evaluate the quality of their synthesis and judgment, not the polish of the AI output
- Teach prompt craft and critical evaluation of AI output as explicit course skills

Further reading: Sweller (1988) Cognitive Load Theory | Bloom et al. (1956) Taxonomy | Lave & Wenger (1991) Situated Learning | Brown et al. (1989) Situated Cognition

Handout organized by Claude