Article

Not peer-reviewed version

AI-Enhanced Teaching: A Strategic Framework for Schools and Colleges

Basker Palaniswamy *

Posted Date: 13 March 2026

doi: 10.20944/preprints202603.1031.v1

Keywords: artificial intelligence in education; AI-assisted teaching; intelligent tutoring systems; educational technology; AI-enhanced learning; digital education; human-AI collaboration in teaching; adaptive learning; classroom innovation; future of education

 doi:10.20944/preprints202603.1031.v1

Article

AI-Enhanced Teaching: A Strategic Framework for Schools and Colleges

Basker Palaniswamy

Insight Research Ireland Centre for Data Analytics, University College Cork, Cork City, Ireland, European Union; basker170889@gmail.com or basker170889@zohomail.eu

Abstract

Artificial intelligence is rapidly transforming education. Tools such as modern AI language models can now generate essays, explain complex concepts, create lesson plans, produce quizzes, and summarize entire textbooks within seconds. For many teachers and institutions, this raises an important question: *what is the role of a human educator in an age when machines can instantly provide information?* This paper presents an accessible framework that helps schools and colleges integrate artificial intelligence into teaching while preserving the essential human elements of education. Rather than viewing AI as a replacement for teachers, the framework positions AI as a powerful assistant that can support lesson preparation, personalized feedback, and adaptive learning resources. By automating repetitive tasks such as content generation, grading support, and material organization, AI allows educators to focus on what machines cannot easily replicate: mentorship, creativity, ethical reasoning, critical thinking, and inspiration. The framework outlines practical strategies for using AI responsibly in classrooms, including guidelines for AI-assisted lesson planning, student engagement techniques, and safeguards to maintain academic integrity. It also discusses how institutions can prepare both teachers and students for an AI-augmented learning environment by promoting digital literacy, responsible tool usage, and critical evaluation of AI-generated information. Ultimately, the goal of AI-enhanced teaching is not to replace educators, but to empower them. When used thoughtfully, artificial intelligence can reduce administrative workload, expand access to high-quality learning resources, and create more personalized educational experiences. In this vision, AI becomes a supportive partner, while teachers remain the guiding force who cultivate curiosity, wisdom, and human understanding in the classroom.

Keywords: artificial intelligence in education; AI-assisted teaching; intelligent tutoring systems; educational technology; AI-enhanced learning; digital education; human-AI collaboration in teaching; adaptive learning; classroom innovation; future of education

---

1. Introduction: The AI Revolution in Education

Empowering Educators to Thrive—Not Just Survive — in the Age of GPTs and Artificial Intelligence. Artificial intelligence has entered the classroom—not as a distant futuristic concept, but as a present, transformative reality. Large language models (GPTs) can now draft essays, solve differential equations, generate code, summarise textbooks, and even produce lesson plans. For educators, this raises a profound question: *What is the role of a teacher when a machine can deliver information instantly?* [1]

The answer is both reassuring and challenging. AI excels at pattern recognition, retrieval, and generation [2]. It does **not** excel at empathy, ethical judgment, cultural sensitivity, creative vision, or the ability to inspire a room of restless teenagers to care about photosynthesis [3]. These remain deeply human capabilities—and they are precisely what great teaching is about.

Consider the scale of the transformation. A single AI model can now generate, in seconds, what once took hours: a complete lesson plan with differentiated activities, a set of twenty exam questions spanning Bloom's taxonomy [4,5], a personalised feedback letter for every student in a class

 doi:10.20944/preprints202603.1031.v1

of forty, or a literature review covering fifty recent papers. The question is not whether AI will change education—it already has. The question is whether educators will *lead* this change or be swept along by it.

Golden Rule

AI delivers information.

YOU deliver **INSIGHT**, **VISION** & **JUDGMENT**.

This article presents two complementary frameworks—one for **schools** (elementary through higher secondary) and one for **colleges** (across six major disciplines: Arts, Science, Engineering, Agriculture, Pharma, and Medical)—that transform AI from a threat into the most powerful pedagogical tool educators have ever had. Each framework is structured around a **one-hour session** divided into four 15-minute phases, designed so that every lesson begins with human curiosity and ends with human ownership.

The frameworks are not theoretical abstractions. They are practical, classroom-tested structures that any educator can adopt immediately, regardless of their technical proficiency with AI. The only prerequisite is a willingness to see AI as a *collaborator in pedagogy*—a tool that handles the mechanical so that the human can focus on the meaningful.

Survivability Principle #1: Reframe the Narrative

AI does not replace teachers. AI replaces the *tasks* teachers used to do manually—grading, content retrieval, example generation [6]. This frees teachers to do what only humans can: mentor, inspire, provoke, and connect [7]. The educators who thrive will be those who embrace this reframing earliest and most completely.

1.1. The Landscape of AI in Education Today

To understand why these frameworks matter, consider the current landscape. As of 2025–2026, the following capabilities are widely available at little or no cost to educators and students:

- **Text generation:** AI can produce essays, reports, summaries, translations, and creative writing in seconds, across virtually any subject and any level of complexity.
- **Code generation:** AI can write, debug, and explain code in dozens of programming languages, from simple scripts to complex algorithms.
- **Image generation:** AI can create illustrations, diagrams, concept art, and photorealistic images from text descriptions.
- **Data analysis:** AI can process datasets, identify patterns, generate visualisations, and produce statistical summaries.
- **Personalised tutoring:** AI can adapt explanations to a student's level, answer follow-up questions, and provide step-by-step guidance through complex problems.
- **Research assistance:** AI can summarise papers, identify relevant literature, and synthesise findings across multiple sources.

These capabilities are remarkable—and they are improving rapidly [2,8]. But they share a common limitation: they operate entirely within the domain of *information processing*. They cannot observe a student's body language, sense the energy of a classroom, make ethical judgments about competing values, or build the trust that enables a struggling teenager to ask for help. These remain the province of human educators, and they are the foundation upon which great teaching has always rested.

The more tasks AI automates, the more valuable the *non-automatable* skills become [9,10]. Paradoxically, the rise of AI makes human skills—empathy, judgment, creativity, mentorship—*more* important, not less. Educators who invest in developing these skills, both in themselves and in their students, are building the most durable competitive advantage possible.

2. Part I: AI-Powered Classroom—Framework for Schools

The school-level framework recognises a fundamental truth: younger learners need **curiosity before technology** [11,12]. A child who has never wondered why the sky is blue will not benefit from an AI that can explain Rayleigh scattering. The framework therefore begins with human wonder and progressively introduces AI as a tool—never as a master.

The framework is designed for a single one-hour session, but its principles can be applied to any lesson, any subject, and any age group. The four phases—Spark & Explore, Create & Collaborate, Apply & Build, and Reflect & Own—form a pedagogical arc that moves from curiosity through critical engagement to creative production and, finally, to reflective ownership. This arc ensures that AI never becomes the centre of the lesson; the *student* always is.

2.1. Phase 1: Spark & Explore (0–15 Minutes)

Phase 1 — Spark & Explore (0–15 min)

Core Activities:

- Teacher poses a real-world **WONDER QUESTION**
- Students brainstorm answers *without devices*
- Teacher demos asking the SAME question to GPT
- Class compares: What WE think vs. what AI said

KEY: Build curiosity BEFORE technology.

The purpose of this phase is to establish that *human thinking comes first*. Before any student sees an AI response, they must have committed to their own ideas. This prevents the most dangerous outcome of AI in classrooms: the death of independent thought before it has a chance to develop.

2.1.1. Elementary Level (Ages 5–11)

At the elementary level, the wonder question must connect to the child's *immediate, sensory world*. Young children think concretely; abstraction comes later. Effective wonder questions for this age include:

- "Why do leaves change colour in autumn?"
- "How does a caterpillar become a butterfly?"
- "What would happen if it never rained?"
- "Why is the ocean salty but rivers are not?"
- "How do birds know where to fly in winter?"

The teacher writes student answers on the board—misspellings and all—because the point is *ownership*. A child who says "the leaf gets sleepy" has engaged their imagination in a way that no AI-generated answer can replicate. The teacher then asks the same question to ChatGPT or Claude, projected on the screen. Students observe the AI's response and the class discusses: *Whose answer was more interesting? Whose made you think more? Did the computer say anything surprising?*

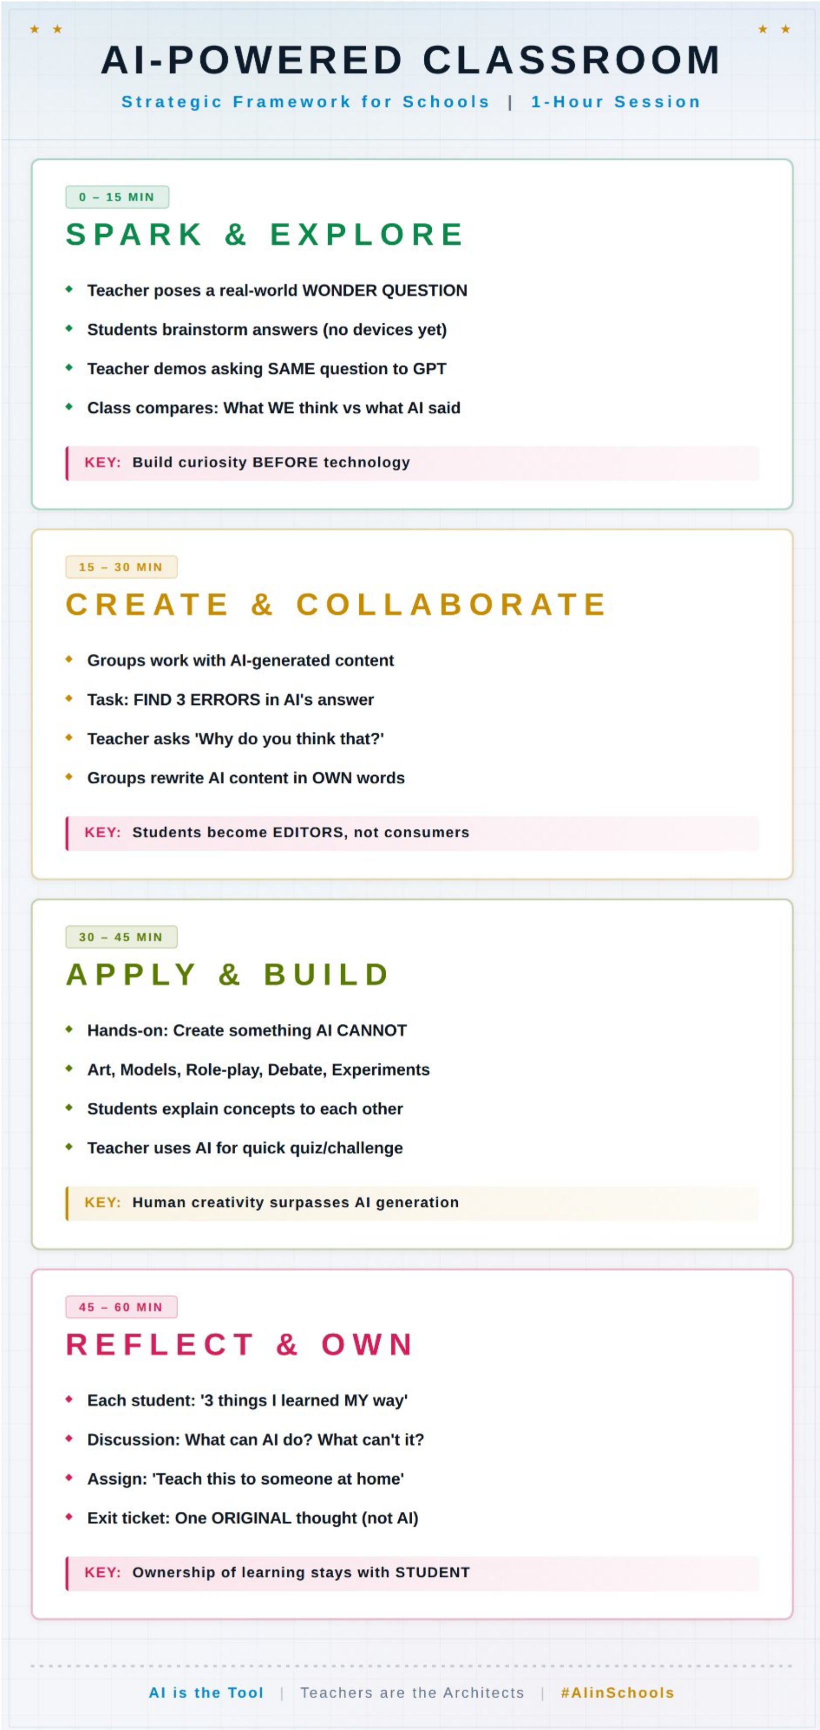

Figure 1. AI-Powered Classroom: Strategic Framework for Schools—a 1-hour session structure progressing from Spark & Explore through Reflect & Own, with key principles anchoring each phase.

At this age, the comparison is not about accuracy—it is about *engagement*. The child who said "the leaf gets sleepy" has expressed a metaphor, a creative leap, a moment of genuine cognitive effort. The AI's technically accurate response about chlorophyll and temperature changes is informative but lacks this spark. The teacher highlights this distinction explicitly: "Your brain did something amazing just now—it made a connection that no computer would have made. That's called *imagination*, and it's something only humans have."

For very young children (ages 5–7), the wonder question should be paired with a physical experience. Before asking "Why do things fall down?", let the children drop different objects and observe. Before asking "Why is ice slippery?", let them hold ice cubes and describe what they feel. The physical experience anchors the question in the body, making it real in a way that a screen never can.

Use "think-pair-share" before the AI demo [13]. Children first think alone (30 seconds), then share with a partner, then share with the class. This ensures every child has already committed to their own idea *before* seeing the AI's answer, preventing passive consumption. For younger children (ages 5–7), use drawing instead of writing: "Draw what you think the answer looks like."

Mrs. O'Sullivan's Year 3 class in Cork was studying magnets. Before showing any AI output, she asked each child to draw what they thought was happening *inside* the magnet. One child drew tiny invisible hands reaching out; another drew swirling circles; a third drew the magnet as a hungry animal "eating" the metal. She then asked Claude the same question and projected the response. The children compared their drawings with the AI's explanation of magnetic fields. The result? Every child remembered the lesson three weeks later—because their own drawing was the anchor. The AI's explanation enriched their understanding; their own creativity made it memorable.

2.1.2. Higher Secondary Level (Ages 14–18)

Higher secondary students are capable of metacognition—thinking about thinking [14]. The wonder question can therefore be more layered, more provocative, and more connected to the complexities of the real world:

- "Is social media making us more connected or more lonely?"
- "Should gene editing be allowed in human embryos?"
- "Can an AI-generated poem be considered real art?"
- "Is economic growth always good for a country?"
- "Should a self-driving car sacrifice its passenger to save five pedestrians?"

Students write a 2-minute *quick-take* paragraph, then the teacher prompts GPT with the exact same question. The class analyses not just the *content* of the AI's response but its *structure*, *tone*, and *assumptions*. Higher secondary students should be encouraged to ask: "What perspective is missing from the AI's answer? Whose voice is not represented? What values does the AI seem to prioritise, and why?"

This phase is particularly powerful for developing *epistemic awareness*—the understanding that all knowledge is produced from a particular standpoint, with particular assumptions [15], and that AI outputs are no exception. Students who develop this awareness become resistant to misinformation, propaganda, and the uncritical consumption of AI-generated content.

At this level, the teacher should also introduce the concept of *prompt sensitivity*: the same question, phrased slightly differently, can produce significantly different AI responses. Demonstrate this by asking the same question three ways—neutral, leading, and provocative—and showing how the AI's

 doi:10.20944/preprints202603.1031.v1

answer shifts. This exercise reveals that AI is not a neutral oracle but a system that responds to the framing of the input, just as human answers are shaped by how questions are asked.

Mr. Kavanagh's Transition Year class debated the trolley problem with a twist: he asked Claude to generate its answer first. The AI produced a balanced analysis of utilitarian vs. deontological perspectives. But when students read it, one asked: "Why does the AI assume the five people are equally valuable? What if one of them is a convicted criminal and the one on the side track is a surgeon who saves lives?" This single question opened a 20-minute discussion about moral complexity that the AI's response—competent but bloodless—could never have generated. The students discovered that ethics requires not just analysis but *moral imagination*, something AI fundamentally lacks.

Survivability Principle #2: Question the Machine

Students who can critically evaluate AI outputs are more valuable in the workforce than students who can merely use AI. Teach interrogation, not just operation. The student who asks "Why did the AI say that?" is developing a skill that will remain valuable long after today's AI models are obsolete.

2.2. Phase 2: Create & Collaborate (15–30 Minutes)

Phase 2 — Create & Collaborate (15–30 min)

Core Activities:

- Groups work with AI-generated content
- Task: FIND 3 ERRORS in AI's answer
- Teacher asks "Why do you think that?"
- Groups rewrite AI content in OWN words

KEY: Students become EDITORS, not consumers.

This phase transforms the student's relationship with AI-generated content from passive consumption to active critical engagement. The core insight is that in a world flooded with AI-generated text, the most valuable skill is not the ability to *produce* content—AI can do that—but the ability to *evaluate* it.

2.2.1. Elementary Level (Ages 5–11)

For young children, "finding errors" is gamified as a treasure hunt. The teacher generates a simple AI passage—say, a paragraph about penguins—and deliberately includes or highlights factual errors ("Penguins live in the Arctic", "Penguins can fly short distances"). Students work in pairs with coloured markers to circle errors. This develops critical reading skills while making AI feel approachable rather than authoritative.

The rewriting exercise at this level is verbal: children *retell* the corrected passage in their own words to their partner. The teacher circulates and asks each pair: "How would *you* explain this to a younger child?" This develops two skills simultaneously: content mastery (they must understand the material to explain it) and communication (they must adapt their language to their audience).

For older elementary students (ages 9–11), the error hunt becomes more sophisticated. The teacher generates a passage that is *mostly* correct but contains subtle errors—not outright falsehoods, but oversimplifications, missing context, or misleading implications. For example, an AI passage about the solar system might say "Pluto is a planet" (contested), or a passage about healthy eating might imply

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

that all fats are bad (oversimplified). These exercises teach children that "correct" and "incorrect" are not always binary categories—a crucial lesson for navigating an information-rich world.

Avoid letting the "error hunt" become a simple true/false exercise. Encourage students to explain *why* something is wrong, not just *that* it is wrong. "Penguins don't fly because..." builds reasoning; a checkmark does not. For younger children, use sentence starters: "I think this is wrong because..." or "A better way to say this would be..."

Ms. Murphy's Year 5 class was studying the Vikings in Ireland. She asked Claude to write a paragraph about Viking life and deliberately selected a response that described Vikings as "savage warriors who only raided and destroyed." Students, who had been studying Viking craftsmanship, trade routes, and settlement patterns, were outraged. "That's not fair!" said one student. "They also built Dublin!" The class rewrote the paragraph collaboratively, producing a nuanced account that acknowledged both raiding and settlement. The AI's oversimplification became the catalyst for the students' most sophisticated historical writing of the year.

2.2.2. Higher Secondary Level (Ages 14–18)

At this level, the errors are not factual blunders but *subtle weaknesses*: logical fallacies, oversimplifications, missing context, cultural bias, lack of citations, or failure to acknowledge uncertainty. The teacher generates an AI essay on a topic the class is studying and distributes it to groups. Each group must:

1. Identify three substantive weaknesses (not typos or formatting issues)
2. Explain *why* each is a weakness, citing specific evidence or reasoning
3. Rewrite the relevant paragraph to fix it, demonstrating the improvement
4. Reflect on what the exercise reveals about AI's limitations in this subject area

This exercise transforms students from passive readers into active critical editors—a skill that directly transfers to university-level work and professional life. Students discover that AI outputs, while impressively fluent, often lack depth, nuance, and the ability to handle genuine complexity.

Higher secondary students should also explore the concept of *AI hallucination*—the phenomenon where AI generates plausible-sounding but entirely fabricated information [16,17]. Provide students with an AI-generated passage that contains a hallucinated citation (a paper that does not exist, attributed to a real researcher) and challenge them to verify it. This exercise is profoundly important for developing research integrity skills.

Survivability Principle #3: The Editor Outlasts the Writer

In a world where AI can generate unlimited text, the human who can evaluate, refine, and judge quality becomes indispensable. Every profession—from journalism to medicine to law—will increasingly require people who can critically assess AI-generated content. Train students to be editors, not just writers.

2.3. *Phase 3: Apply & Build (30–45 Minutes)*

Phase 3 — Apply & Build (30–45 min)

Core Activities:

- Hands-on: Create something AI CANNOT
- Art, Models, Role-play, Debate, Experiments
- Students explain concepts to each other
- Teacher uses AI for quick quiz/challenge

KEY: Human creativity surpasses AI generation.

This is the phase where learning becomes *embodied*—where knowledge moves from the head into the hands, the voice, the body. AI can process information and generate text, but it cannot build, perform, experiment, debate, or create physical artefacts. This phase deliberately occupies the space that AI cannot reach.

2.3.1. Elementary Level (Ages 5–11)

AI cannot build a volcano out of papier-mâché, perform a puppet show about the water cycle, or draw a picture that expresses how a child *feels* about the rainforest. Activities at this level should be multi-sensory, collaborative, and joyful:

- **Science:** Build a simple circuit, grow bean plants, conduct sink-or-float experiments, create a weather station from household materials
- **Art:** Create a collage, paint a mural, design a poster using only hands and materials, sculpt with clay or playdough
- **Language:** Role-play a story scene, perform a class debate on a simple motion, create a radio play with sound effects
- **Maths:** Use physical manipulatives (blocks, beads, fraction tiles) to solve problems the AI solved abstractly, measure real objects and compare with AI estimates
- **Geography:** Build a 3D map of the local area, conduct a mini-survey of classmates and graph the results by hand

The teacher can use AI in the background during this phase—generating a quick quiz on the topic, creating differentiated worksheets for different ability levels, or producing a fun challenge question for early finishers. But the *students* are doing hands-on work. The AI is the teacher's assistant, not the student's replacement.

After the build activity, pair students and ask each to explain what they made and what they learned to their partner. This serves as both formative assessment (the teacher can listen for misconceptions) and consolidation (the act of explaining deepens understanding). For younger children, use the prompt: "Pretend your partner is an alien who has never seen Earth. Explain this to them."

2.3.2. Higher Secondary Level (Ages 14–18)

Higher secondary students should tackle projects that require *synthesis, judgment, and originality*—qualities that cannot be automated:

- **Science:** Design and run an actual experiment, then compare results to AI predictions. Where did reality diverge from the model? What variables did the AI not account for?
- **History:** Stage a mock trial of a historical figure, using AI-generated evidence that students must cross-examine for bias, anachronism, and omission

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

- **Literature:** Write a creative response to a poem—not an analysis, but a *conversation* with the text
- **Business Studies:** Develop a startup pitch for a *local* problem, using AI for market research but human judgment for the value proposition
- **Mathematics:** Model a real-world scenario using both hand calculations and AI, then compare approaches and identify where human intuition adds value
- **Languages:** Translate a passage using AI, then improve the translation by adding cultural nuance and idiomatic expressions that the AI missed

The "AI Can't Do This" Test

Before assigning any project, apply this test: *Could a student complete this entirely by pasting the prompt into ChatGPT?* If yes, redesign the assignment. Assignments that require physical creation, personal experience, ethical judgment, local context, real data collection, or live performance are naturally AI-proof. The goal is not to *prevent* AI use, but to design tasks where AI use alone produces an inferior result.

2.4. Phase 4: Reflect & Own (45–60 Minutes)

Phase 4 — Reflect & Own (45–60 min)

Core Activities:

- Each student: "3 things I learned MY way"
- Discussion: What can AI do? What can't it?
- Assign: "Teach this to someone at home"
- Exit ticket: One ORIGINAL thought (not AI)

KEY: Ownership of learning stays with STUDENT.

Reflection is where learning becomes permanent. Without this phase, the lesson remains an experience; with it, the lesson becomes *knowledge*—integrated into the student's understanding of both the subject and of themselves as a learner.

2.4.1. Elementary Level (Ages 5–11)

Young children reflect through *doing and sharing*. The "teach someone at home" assignment is particularly powerful: a child who can explain what they learned to a parent or sibling has truly internalised the knowledge. The exit ticket at this level can be a drawing, a single sentence, or even a verbal statement recorded by the teacher.

Effective reflection prompts for this age group include: "What was the most surprising thing you learned today?" "What did you learn today that a computer couldn't have taught you?" "If you could ask one more question about today's topic, what would it be?" These prompts develop metacognitive awareness—the ability to think about one's own thinking—which research consistently identifies as one of the most powerful predictors of academic success [18,19].

"Draw one thing you learned today that a computer *couldn't* have taught you. Write one sentence about why." This simple exercise builds metacognitive awareness from an early age and reinforces the message that human learning involves more than information retrieval.

For older elementary students, introduce the concept of a *learning journal*. Each week, students spend five minutes writing about what they learned, how they learned it, and what questions they still have. Over time, this journal becomes a powerful record of intellectual growth—and a concrete reminder that learning is a human process, not a mechanical one.

doi:10.20944/preprints202603.1031.v1

2.4.2. Higher Secondary Level (Ages 14–18)

At this level, reflection should be *philosophical and self-aware*. Students discuss:

- What did AI help me understand better today?
- Where did AI mislead me or give me a shallow answer?
- What can I do that AI cannot, and how do I develop that further?
- If I had to explain today's topic to someone with no internet access, how would I do it?
- What assumptions did I make before I started, and how have they changed?
- What would I want to investigate further, and why?

The exit ticket at this level is a written paragraph containing *one original thought*—an insight, question, or connection that the student generated themselves, not prompted by or derived from AI. This trains the habit of original thinking in a world saturated with generated content.

In Ms. Brennan's Leaving Cert Biology class, students studied enzyme kinetics using both traditional methods and AI-generated explanations. For the exit ticket, one student wrote: "I wonder if enzymes in organisms living near deep-sea vents work differently because of the extreme pressure, and whether our models of enzyme kinetics would need to be completely rewritten for those conditions." This was not in any textbook or AI output—it was a genuine moment of scientific curiosity, triggered by the lesson but produced entirely by the student's own thinking. Ms. Brennan later used this question as the starting point for the next week's lesson on extremophile biochemistry.

Survivability Principle #4: Originality Is the New Literacy

In the age of AI, the ability to produce an original thought—not just retrieve or rephrase information—is the defining skill of an educated person. Every lesson should end with students practising this skill. Originality is not a talent possessed by the few; it is a habit developed by the many, through consistent practice and encouragement.

3. Part II: AI-Enhanced Teaching—Framework for Colleges

The college-level framework operates on a more sophisticated premise: university students are training to become *professionals* in specific disciplines. AI's role, limitations, and ethical implications differ dramatically between, say, a pharmaceutical researcher and a studio artist. A pharmacist who trusts an AI hallucination could prescribe a fatal drug interaction; an artist who delegates their creative vision to AI ceases to be an artist. The stakes, the skills, and the safeguards are discipline-specific—and the framework must be too.

The framework provides discipline-specific guidance across six fields: **Arts**, **Science**, **Engineering**, **Agriculture**, **Pharma**, and **Medical**. Each discipline is treated as a *distinct culture of knowledge*—with its own epistemology (how knowledge is produced) [20], its own ethics (what counts as responsible practice), and its own relationship with AI (where the tool helps and where it hinders).

 doi:10.20944/preprints202603.1031.v1

Figure 2. AI-Enhanced Teaching: The comprehensive college-level framework across six disciplines—Arts, Science, Engineering, Agriculture, Pharma, and Medical—structured in four 15-minute phases with discipline-specific activities, culminating in the Golden Rule.

doi:10.20944/preprints202603.1031.v1

3.1. Phase 1: Engage & Provoke (0–15 Minutes)

The first phase challenges students to *confront* AI-generated output with disciplinary scepticism. The AI produces something; the student evaluates it through the lens of their field's standards, ethics, and epistemology. The goal is **productive scepticism**—not rejection of AI, but rigorous evaluation that sharpens professional judgment.

Phase 1 — Engage & Provoke (0–15 min)

AI generates a disciplinary output. Students critique it using field-specific standards. Each discipline brings its own lens: the artist asks about soul, the scientist about testability, the engineer about tradeoffs, the farmer about local conditions, the pharmacist about precision, and the doctor about the patient behind the data.

3.1.1. Arts: "Is This REAL Art?"

Arts — Engage & Provoke

Activities: Ask GPT to write a poem/essay. Critique: Is this REAL art? What's missing—soul, context, lived experience? Introduce the HUMAN element. Peer feedback: originality vs. technical competence.

The arts represent perhaps the most philosophically provocative case for AI in education [21,22]. When a machine can generate a sonnet in the style of Shakespeare, paint in the manner of Monet, or compose music that sounds like Bach, the fundamental question becomes: *What is art, and what makes it valuable?*

The provocation here is not technical but existential. Students compare an AI-generated poem with a human poem on the same theme. The discussion centres on intentionality (did the AI *mean* to write this?), vulnerability (great art often comes from personal risk), context (a poem about grief by someone who has lost a parent carries weight the same words from a statistical model do not), and authorship (if a student curates AI output, who is the artist?).

Professor Walsh's creative writing seminar asked GPT-4 to write a love poem. The result was technically accomplished—good meter, vivid imagery, elegant structure. But when students compared it with a poem written by a classmate about their grandmother's hands, the difference was palpable. "The AI poem sounds like love," one student said. "The other poem *is* love." This distinction—between simulation and expression of emotion—became the central theme of the semester.

If a student submits AI-generated work as their own, is this plagiarism—or collaboration? If an artist uses AI to generate initial concepts and then refines them by hand, where does the AI's contribution end and the artist's begin? These are the central debates of contemporary art practice, and students should engage with them directly.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

3.1.2. Science: “Science Needs DOUBT”

Science — Engage & Provoke

Activities: GPT generates a hypothesis. Evaluate: Is this testable? What assumptions did AI make? Find assumptions based on training data patterns. Key principle: “Science needs DOUBT.”

Science is built on scepticism. Every hypothesis must be testable, every claim falsifiable, every result reproducible [23]. AI generates outputs based on statistical patterns—it does not *understand* the scientific method [24] and cannot distinguish between a plausible-sounding hypothesis and a genuinely testable one.

Students evaluate the AI hypothesis for testability, falsifiability, implicit assumptions, edge cases, and novelty. They ask: What model is the AI relying on? What data informed this? Under what conditions would this fail?

The AI hypothesises nitrogen runoff as the primary cause of algal blooms in a local lake. Students discover the AI missed the new housing development increasing phosphorus loading, the unusually warm spring raising water temperatures, and the decline in a key grazer species. The AI’s hypothesis was a reasonable starting point, but real science requires *local observation, measurement, and judgment*.

3.1.3. Engineering: “Engineering = Tradeoffs”

Engineering — Engage & Provoke

Activities: GPT produces a design solution. Find 3 practical flaws. Identify constraints AI overlooked. Key principle: “Engineering = tradeoffs.”

Engineering is not about finding the *best* solution; it is about finding the *best achievable* solution given competing constraints of safety, cost, time, materials, environmental impact, regulation, and human factors. AI tends to optimise for a single objective without accounting for the messy reality of manufacturing tolerances, supply chain disruptions, local building codes, or changing client requirements.

Students evaluate the AI design as if conducting a formal design review: practical buildability, safety and failure modes, economic viability, ethical implications, and human factors including accessibility, ergonomics, and user experience.

The AI designed a cable-stayed pedestrian bridge for Cork. Students found three critical issues: (1) the design assumed uniform soil conditions, but the site has variable bedrock depth; (2) cable anchorage didn’t account for Ireland’s prevailing southwest winds; (3) the aesthetic design created maintenance access problems tripling lifecycle cost. Students redesigned with a less dramatic but far more practical solution, learning that engineering judgment cannot be automated.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

3.1.4. Agriculture: "Farming = Reading the Land"

Agriculture — Engage & Provoke

Activities: GPT suggests crop rotation plan. Evaluate for local soil/climate/seasonal conditions. Identify what AI missed: microclimate, water table, pests, traditions. Key principle: "Farming = reading the land."

Agriculture is perhaps the most *place-based* of all disciplines [25]. A crop rotation plan that works brilliantly in Iowa's black soils may be catastrophic in the thin, acidic soils of West Cork. AI, trained on global data, produces recommendations that are statistically average—and therefore wrong for any specific field.

Students evaluate the AI plan against soil type, pH, drainage, organic matter, local rainfall patterns, frost dates, biodiversity implications, economic viability, and indigenous knowledge. The farmer's observations about drainage patterns, traditional planting calendars, and local pest behaviour are knowledge forms that no dataset captures.

Students received an AI rotation plan for a West Cork dairy farm recommending maize as a summer crop. Local farmers immediately identified the problem: West Cork's cool, wet summers rarely provide enough growing degree days for maize maturity. Students replaced maize with a multi-species sward (grass, clover, plantain, chicory) that local farmers had been developing for years. AI can process vast datasets, but it cannot *read the land* the way a farmer who has walked it for thirty years can.

3.1.5. Pharma: "Pharma = Precision, Not Speed"

Pharma — Engage & Provoke

Activities: GPT drafts drug interaction report. Verify against pharmacopoeia and official databases. Find dosage/safety gaps; identify AI hallucinations. Key principle: "Pharma = precision, not speed."

Pharmaceutical science operates in a domain where errors are potentially fatal [26]. A drug interaction missed by AI, a dosage recommendation based on hallucinated data, or an unflagged contraindication could result in patient harm or death. This makes critical evaluation of AI outputs a matter of professional ethics and patient safety.

Students verify every claim for accuracy against official databases, completeness (including food and supplement interactions), hallucination detection, patient-specific factors (age, weight, renal/hepatic function, genetic polymorphisms), and jurisdictional regulatory status.

A 72-year-old is prescribed warfarin, amiodarone, and omeprazole. The AI correctly identifies the warfarin–amiodarone interaction (increased bleeding risk) but misses the omeprazole–warfarin interaction (altered CYP2C19 metabolism) and hallucinates a non-existent amiodarone–calcium interaction. Cross-referencing with the BNF, students discover that trusting this report without verification could lead to either insufficient anticoagulation monitoring or unnecessary medication changes. *In pharmacy, verification is not optional—it is the entire point.*

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

Unlike most fields where AI hallucination is merely embarrassing, in pharmaceutical science it is *dangerous*. A fabricated drug interaction, a hallucinated dosage range, or an invented contraindication can directly harm patients. Students must develop the reflex of *systematic verification*—never accepting an AI pharmaceutical claim without independent confirmation from an authoritative source.

3.1.6. Medical: "Medicine = Human Observation"

Medical — Engage & Provoke

Activities: GPT generates differential diagnosis from symptom list. Check against patient history and clinical context. Identify context AI cannot access: affect, body language, family dynamics. Key principle: "Medicine = human observation."

Medicine is fundamentally a *human* discipline—not because its knowledge base is non-computational, but because its practice requires the integration of scientific knowledge with empathy, observation, communication, and ethical judgment. A machine can process lab values faster than any human; but a machine cannot notice the patient's anxiety, their reluctance to make eye contact, or the family dynamics in the room.

Students identify what additional information a clinician would gather beyond the AI's differential: the patient's narrative [27], social determinants of health [28], cultural context, non-verbal cues, and the therapeutic relationship itself.

Medical students were presented with a 45-year-old male complaining of persistent fatigue, weight gain, and difficulty concentrating. The AI generated a differential focusing on thyroid dysfunction, sleep apnoea, and anaemia. When students conducted the simulated consultation, the "patient" revealed, only after gentle, persistent questioning, that he had been feeling hopeless since his wife's death eight months ago. The diagnosis was major depressive disorder—which the AI missed entirely because "sadness" wasn't on the symptom list. Medicine requires the ability to *see what isn't on the chart*.

3.2. *Phase 2: Deepen & Challenge (15–30 Minutes)*

In this phase, students move from evaluation to *creation*—but with AI as a brainstorming partner, not an author. The student's own vision, methodology, and judgment take centre stage.

Phase 2 — Deepen & Challenge (15–30 min)

Students create their OWN version. AI serves as a brainstorming partner only. The student leads; the AI assists. Professional judgment, not computational power, drives every decision.

3.2.1. Arts: Create YOUR Vision

Arts — Deepen & Challenge

Students create their OWN version of the AI's work. AI generates possibilities; students curate with personal vision. They compare "AI vs. your vision" and receive peer feedback focused on originality, not polish. The emphasis shifts to *voice*—the student's unique creative identity. Students

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

also experiment with AI as material: just as a collage artist uses found images, a contemporary artist might reshape AI output through selection, arrangement, and transformation.

Create three works: one AI-generated, one collaborative with AI, one entirely human-made. For each, write a 200-word artist's statement explaining the process, the decisions, and what is gained or lost. This develops critical self-awareness about technology's role in creative practice.

3.2.2. Science: Design YOUR Experiment

Science — Deepen & Challenge

Design an experiment to test the AI's hypothesis. AI finds papers and data; students identify methodology gaps and design controls. The key insight: experimental design is a *creative act* requiring scientific imagination—the ability to envision what might go wrong, what confounding variables lurk, and what results would mean in broader context. Students also use AI to stress-test their design: ask the AI for weaknesses, then evaluate which critiques are valid.

The AI suggested a novel metal oxide catalyst for hydrogen-producing water splitting. Students must design the experiment: substrate concentrations, temperatures, pressures, controls, measurement methods. They discover the AI assumed access to a gas chromatograph their lab doesn't have. They redesign using available instruments, learning that real science operates within constraints that algorithms don't face.

3.2.3. Engineering: Redesign with Real Constraints

Engineering — Deepen & Challenge

Redesign with real constraints: material costs, safety codes, manufacturing tolerances. AI assists with calculations and code snippets; the human optimises the cost/safety/ethics tradeoff. Critical requirement: prototype with pen and paper *first*. Students also practise *design justification*—explaining not just *what* they designed but *why* they made each choice. In professional engineering, defending a decision under questioning is as important as producing the design.

If an AI-assisted design fails, who bears responsibility? The answer, in every jurisdiction, is clear: the licensed professional engineer is responsible [29]. AI assistance does not transfer liability. Students must grapple with this reality before entering a profession where their signature on a drawing carries legal weight.

3.2.4. Agriculture: Design YOUR Field Trial

Agriculture — Deepen & Challenge

Design a field trial protocol. AI reviews soil/weather data; students add indigenous wisdom, local farmer knowledge, and sustainable practices. This phase emphasises *knowledge integration*—bringing together quantitative data (AI, sensors) with qualitative knowledge (farmers, communities, ecosystems). An AI recommending maximum nitrogen may optimise yield; a farmer who applies less knows the excess will damage the downstream fishery.

Where possible, pair this phase with an actual farm visit. Have students walk the fields they are designing for, speak with the farmer, and observe the landscape with their own eyes. The gap between the AI's satellite-derived recommendation and the reality on the ground is often startling—and deeply instructive.

3.2.5. Pharma: Design YOUR Drug Trial

Pharma — Deepen & Challenge

Design clinical trial methodology. AI cross-references contraindications and calculates statistical power; students assess ethical risks and regulatory compliance. The *ethical* decisions—inclusion criteria, informed consent, adverse event monitoring, and whether the trial is justifiable at all—must be made by humans. Students check designs against ICH-GCP guidelines [30]. AI retrieves the guidelines; students interpret and apply them.

Pharmaceutical research operates under *primum non nocere*—first, do no harm [31]. AI can accelerate drug discovery [32], but the decision to test a compound on a human being is a moral decision, not a computational one. No amount of AI confidence substitutes for rigorous clinical testing and ethical review.

3.2.6. Medical: Build YOUR Clinical Reasoning

Medical — Deepen & Challenge

Build a clinical reasoning pathway for a complex case. AI summarises case studies and provides evidence references; students add patient empathy and navigate ethical dilemmas. The pathway is not a flowchart—it is a narrative: why consider this diagnosis? What ruled it out? What additional tests? How to communicate it? What if the patient refuses treatment? This phase also introduces ethical dilemmas: the patient wanting unsupported treatment, the family wanting "everything done" for a dying relative, the resource allocation decision.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

A 6-year-old with a newly diagnosed chronic condition. The AI provides the evidence-based protocol. But the parents disagree: the mother wants to follow the medical recommendation; the father prefers alternative medicine first. Students navigate this, balancing parental autonomy, the child's best interests, and the therapeutic relationship. The "right answer" depends not just on evidence but on values, cultural context, and the clinician's ability to build trust. AI provides evidence; only a doctor can have *the conversation*.

3.3. Phase 3: Build & Synthesize (30–45 Minutes)

The production phase, where disciplinary knowledge, AI assistance, and human judgment converge into a tangible, assessable output.

Phase 3 — Build & Synthesize (30–45 min)

Produce a tangible, discipline-specific output. AI assists with sub-tasks; the student owns the synthesis and the final product. The output must demonstrate both disciplinary competence and critical engagement with AI.

3.3.1. Arts: The Human + AI Mashup

Arts — Build & Synthesize

Create a Human + AI mashup project. Curate, edit, and add meaning. Analyse AI's limitations explicitly and produce a *personal artistic statement*. The final work must include something AI cannot generate: personal narrative, emotional truth, embodied experience, or site-specific context.
Example Projects: A video essay juxtaposing AI-generated imagery with personal footage; a poetry collection where each AI poem is paired with the student's "response" poem; a musical composition using AI-generated harmonic progressions as a foundation for human improvisation; a visual art piece where AI-generated elements are physically painted over, cut up, or transformed by hand.

3.3.2. Science: Run the Experiment

Science — Build & Synthesize

Run the experiment or simulation. AI visualises data patterns, but *YOU interpret* (not AI). Cross-reference with published research, identifying where findings confirm, extend, or contradict existing knowledge.
Key Requirement: Produce a written interpretation that goes beyond what AI could generate from data alone: contextualise within the broader field, identify design limitations, propose follow-up experiments, and speculate (with appropriate caveats) about implications. The scientist's job is not to process data—AI does that—but to *understand what the data means*.

3.3.3. Engineering: Build the Prototype

Engineering — Build & Synthesize

Build a working model, simulation, or prototype. AI assists debugging only—design decisions are frozen from Phase 2. Present design justifications in a formal peer review, defending choices under questioning.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

Key Requirement: Document not only what you built but what you *chose not to build*—alternative designs considered and rejected, and why. This documentation of the decision-making process is often more valuable than the product itself, because it demonstrates professional engineering judgment that no AI can replicate.

3.3.4. Agriculture: The Integrated Farm Plan

Agriculture — Build & Synthesize

Create an integrated farm plan synthesising AI-derived data (yield predictions, weather models, market analysis) with field observations, farmer interviews, and sustainability assessments. AI models yield predictions; students validate with fieldwork and produce a sustainability impact report covering environmental, economic, and social dimensions.
Key Requirement: Include a section on "What the AI Cannot Know"—specific local factors (soil variability within the field, microclimatic effects of hedgerows, farmer knowledge about historically wet corners) that AI models cannot capture. This demonstrates the irreplaceable value of ground-truth observation.

3.3.5. Pharma: Compound Analysis and Safety

Pharma — Build & Synthesize

Formulate a comprehensive compound analysis. AI assists molecular modelling, ADMET prediction, and literature retrieval; students verify lab results, check against official databases, and complete safety and efficacy assessment. Every AI-generated claim must be traceable to verified, peer-reviewed data.
Key Requirement: Produce a "Verification Report" documenting, for each AI claim, the source used to verify it, whether verification confirmed or contradicted the AI, and any discrepancies found. This develops the habit of *systematic verification*—the pharmacist's most important professional skill.

3.3.6. Medical: The Patient Consultation

Medical — Build & Synthesize

Simulate a complete patient consultation: history-taking, examination, differential diagnosis, investigation, management plan, and communication. AI provides clinical references and evidence summaries; students practise bedside manner, shared decision-making, and communicating complex information in language the patient can understand.
Key Requirement: Develop a treatment plan *with rationale*—explaining not just *what* to prescribe, but *why*, and how to communicate it to the patient. Document how you would handle the patient's likely questions, concerns, and emotional responses. This is the dimension of medicine that AI cannot reach: the human encounter.

3.4. Phase 4: Reflect & Lead (45–60 Minutes)

The final phase asks students to step back and think about the *role of AI in their profession*. This is a professional, ethical, and philosophical exercise.

Phase 4 — Reflect & Lead (45–60 min)

Present work, evaluate AI's contribution, and articulate the professional's irreplaceable role. Assignments require original thought that no textbook (or AI) covers.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 13 March 2026 doi:10.20944/preprints202603.1031.v1

3.4.1. Arts: "What I Made AI Can't"

Arts — Reflect & Lead

Reflection: Present work and articulate what makes it distinctly human. Journal on creative identity: "How do I create differently from a machine? What is my voice, my perspective, my vulnerability?" Engage with the ethics of AI in creativity—authorship, originality, the commodification of artistic labour.
Assignment: *Create, not consume.* Produce an original work rooted in personal experience, site-specific context, or embodied performance—and write a 500-word statement defending its human origin. The statement should explain: What could AI *not* have contributed to this work? What makes it irreducibly yours?

3.4.2. Science: "Where AI Helped vs. Where It Misled"

Science — Reflect & Lead

Reflection: Present results with uncertainty analysis. Explicitly document where AI helped (literature review, data visualisation, statistical calculations) and where it misled (incorrect assumptions, hallucinated citations, oversimplified models). Formulate the NEXT research question—one this experiment revealed but could not answer.
Assignment: *Novel, not textbook.* Propose a research question extending beyond existing literature—a question that could not be AI-generated because it requires a creative leap, field observation, or cross-disciplinary connection no dataset contains.

3.4.3. Engineering: "AI as Tool vs. AI as Crutch"

Engineering — Reflect & Lead

Reflection: Demo and defend design decisions formally. Distinguish YOUR work from AI's work. Discuss using AI as a tool (calculations, debugging) versus as a crutch (letting it make design decisions without verification). Document YOUR work vs. AI's work explicitly.
Assignment: *Innovate beyond AI.* Propose a solution to a local engineering problem requiring site-specific knowledge, stakeholder consultation, and professional judgment AI cannot access. Document the site visit, stakeholder conversations, and design rationale.

3.4.4. Agriculture: "AI vs. Farmer's Intuition"

Agriculture — Reflect & Lead

Reflection: Present farm-to-table strategy. Stage a debate: AI recommendations vs. farmer's intuition—which is more reliable, under what conditions? Document field observations AI could not predict. Discuss traditional ecological knowledge in a data-driven world.
Assignment: *Solve a LOCAL problem.* Identify a specific agricultural challenge in your community—soil degradation, water scarcity, market access, biodiversity loss—and develop a solution integrating AI data with local knowledge. Interview at least one local farmer as part of the research.

doi:10.20944/preprints202603.1031.v1

3.4.5. Pharma: "AI Accuracy vs. Expert Review"

Pharma — Reflect & Lead

Reflection: Present drug safety report. Compare AI accuracy against expert pharmacist review. Document specific clinical judgment calls—decisions requiring weighing competing risks, interpreting ambiguous data, or recommending under uncertainty.
Assignment: *Novel formulation idea.* Propose a new approach to drug delivery, formulation, or repurposing inspired by AI-generated leads but going beyond them. Explain how AI contributed, what it could not generate alone, and what further research is needed.

3.4.6. Medical: "The Case NO Textbook Covers"

Medical — Reflect & Lead

Reflection: Present diagnosis and care plan. Compare AI diagnostic suggestions with clinician judgment: where did AI help narrow the differential? Where did it miss something only human observation could catch? Document the patient interaction experience—what you learned that was not in the chart.
Assignment: *The case no textbook covers.* Work through a complex, ambiguous scenario: atypical presentation, multiple comorbidities, conflicting preferences, social circumstances complicating management. Write up the case including reasoning, uncertainties, and how you would handle the conversation with the patient and their family.

4. How Teachers Can Use GPTs to Enrich Teaching

Beyond the classroom session itself, AI offers teachers ten powerful capabilities for *preparation, assessment*, and *personalisation*. These are available today, and educators who adopt them gain significant advantages in efficiency, creativity, and responsiveness to student needs.

10 Ways Teachers Can Use GPTs

1. **Lesson Planning:** Generate diverse examples, analogies, and case studies instantly. A single prompt can produce ten worked examples at different difficulty levels, saving hours of preparation while ensuring variety.
2. **Assessment:** Create differentiated quizzes, rubrics, and question banks at all levels—from recall to synthesis—in minutes rather than hours. Target specific learning outcomes and Bloom's taxonomy levels.
3. **Personalisation:** Present the same concept at beginner, intermediate, and advanced levels. Enable true differentiated instruction at scale.
4. **Smart Feedback:** Draft detailed feedback templates that students then personalise. AI handles the scaffold; the teacher adds the human touch—specific praise, gentle challenge, knowledge of what this student needs.
5. **Research Aid:** Summarise papers, find cross-disciplinary connections fast. Particularly valuable for teachers who teach across multiple subjects.
6. **Debate Partner:** AI argues the opposite position, forcing students to defend their thinking with evidence and logic.
7. **Lab & Field Prep:** Generate starter protocols with deliberate errors for students to find and fix—building scientific rigour through error detection.

8. **Clinical Simulation:** Generate patient scenarios for medical and pharma training, with varying complexity and ambiguity.
9. **Accessibility:** Translate, simplify, and create visual aids for diverse learners, including multilingual students and those with learning differences.
10. **Ethical Reasoning:** AI presents dilemmas; students develop moral frameworks and defend their positions in structured discussions.

For teachers new to AI, start small. Before each lesson, spend five minutes asking an AI: "Give me three unexpected analogies for [today's topic]." Use the best one in class. Within a week, you will have developed an intuition for how AI can enhance your teaching—and where it falls short.

5. Surviving and Thriving in the Age of GPTs

The rise of AI in education is not a crisis—it is a *selection pressure*. Like every transformative technology before it—the printing press, the calculator, the internet [33]—AI will reshape what it means to be educated, what it means to be a professional, and what it means to teach. Educators and students who adapt will not merely survive but will become more effective, more creative, and more irreplaceable than ever before.

5.1. For Teachers: Your Irreplaceable Value

For Teachers: Why You Cannot Be Replaced

- **Relationship:** AI cannot build trust with a struggling student. The student who comes to your office not because they need help with the homework but because they need someone to believe in them—that is your domain, and it always will be.
- **Contextual judgment:** You know that Riya learns best through stories, that Sean shuts down when put on the spot, that the class is exhausted after PE. AI knows none of this.
- **Ethical modelling:** When you say "I don't know, but let's find out together," you are teaching intellectual humility. AI generates confidence; you demonstrate character.
- **Inspiration:** The teacher who made you love a subject did it by caring, by being passionate, by seeing you as an individual. AI can inform; only a human can inspire [7].
- **Curriculum judgment:** You decide what to teach, when, and why. You know this class needs extra time on fractions, this student is ready for extension, this topic connects to something discussed three months ago.
- **Community building:** A classroom is a community where young people learn to collaborate, disagree respectfully, celebrate success, and support each other. You are the architect of that community.

5.2. For Students: Skills That AI Cannot Replace

For Students: Your Competitive Advantage in an AI World

- **Critical thinking:** The ability to evaluate, question, and judge—not just retrieve and reproduce [34]. The student who asks "Is this true?" before "Is this on the exam?" will thrive.

- **Creativity with purpose:** AI generates; humans create with intention, emotion, and meaning. Seeing a problem differently, making unexpected connections, imagining what does not yet exist—this is your superpower.
- **Embodied skills:** Surgery, farming, art-making, engineering inspection, patient examination, lab technique, fieldwork—these require hands, eyes, and physical presence.
- **Ethical reasoning:** Deciding what *should* be done, not just what *can* be done. As AI makes more things possible, reasoning about what is right becomes more important.
- **Communication and leadership:** Persuading a room, calming a patient, mentoring a colleague, building a team, negotiating a conflict—irreducibly human skills.
- **Local and tacit knowledge:** Understanding your community, your soil, your patients, your market—knowledge that cannot be scraped from the internet because it lives in relationships and place.
- **Adaptability:** The AI tools of 2026 will differ from those of 2030. The student who learns *how to learn*—the meta-skill of adapting to new tools and challenges [35]—will always be valuable.

5.3. For Institutions: Strategic Imperatives

🏛 What Schools and Universities Must Do Now

- **Update assessment:** If an assignment can be completed by AI, it is testing the wrong skills [36]. Redesign assessments to require personal reflection, physical creation, oral defence, or local investigation.
- **Train teachers:** Focus professional development not on *how to use AI tools* (this changes too fast) but on *how to teach in an AI-rich environment* (the enduring challenge).
- **Establish AI ethics policy:** Develop clear, nuanced policy distinguishing legitimate AI assistance from academic dishonesty, in consultation with students.
- **Invest in the human:** Double down on mentorship, pastoral care, small-group teaching, and experiential learning—educational experiences AI makes *more* valuable, not less.
- **Embrace transparency:** Teach students to declare AI use, just as they declare sources. The goal is transparent, accountable professional practice.

The students and teachers most at risk are not those who lack AI skills. They are those who use AI *passively*—copying outputs, accepting answers uncritically, outsourcing their thinking. The frameworks in this article are designed to prevent this by structuring every interaction around human agency, critical evaluation, and original thought. Passivity is the enemy; curiosity is the cure.

6. Conclusion: The Human at the Centre

The AI revolution in education is not about technology—it is about *pedagogy*. The question is not "How do we teach students to use AI?" but "How do we teach students to think, create, and lead in a world where AI handles the routine?"

The frameworks presented here—for schools and for colleges—share a common architecture. They begin with human curiosity, not AI output. They engage AI critically, as an object of scrutiny rather than an oracle. They challenge students to build something that AI cannot replicate. And they end with reflection, ensuring that learning stays with the learner.

 doi:10.20944/preprints202603.1031.v1

This four-phase structure ensures that AI is always a *tool in service of human learning*, never a replacement for it. The six discipline-specific expansions—Arts, Science, Engineering, Agriculture, Pharma, and Medical—demonstrate that this principle is not abstract but practical, adapting to the unique epistemology, ethics, and professional demands of each field.

The future of education belongs to those who understand that the most powerful technology in any classroom has always been the same: a curious mind, guided by a caring teacher, engaged with the world in all its complexity and beauty [37]. AI amplifies this; it does not replace it.

★ Golden Rule

AI is the Tool. Teachers are the Architects.

The future belongs to educators who use AI to amplify human potential—not replace it.

Empower. Inspire. Lead.

The classroom is yours. The technology is ready. The students are waiting.

Acknowledgments: This publication has emanated from research supported by a grant from Research Ireland under Grant number 12-RC-2289-P2 which is co-funded under the European Regional Development Fund. For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission.

References

1. Seldon, A.; Abidoye, O. *The Fourth Education Revolution Reconsidered: Will Artificial Intelligence Enrich or Diminish Teaching and Learning?*; University of Buckingham Press, 2020.
2. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arber, S.; von Arx, S.; et al. On the Opportunities and Risks of Foundation Models. *arXiv preprint arXiv:2108.07258* **2021**.
3. Luckin, R.; Holmes, W.; Griffiths, M.; Forcier, L.B. *Intelligence Unleashed: An Argument for AI in Education*; Pearson Education, London, 2016.
4. Bloom, B.S.; Engelhart, M.D.; Furst, E.J.; Hill, W.H.; Krathwohl, D.R. *Taxonomy of Educational Objectives: The Classification of Educational Goals. Handbook I: Cognitive Domain*; David McKay Company, New York, 1956.
5. Krathwohl, D.R. A Revision of Bloom's Taxonomy: An Overview. *Theory into Practice* **2002**, *41*, 212–218.
6. Holmes, W.; Bialik, M.; Fadel, C. *Artificial Intelligence in Education: Promises and Implications for Teaching and Learning*; Center for Curriculum Redesign, Boston, MA, 2019.
7. Hattie, J. *Visible Learning for Teachers: Maximizing Impact on Learning*; Routledge, London, 2012.
8. Kasneci, E.; Sessler, K.; Küchemann, S.; Bannert, M.; Dementieva, D.; Fischer, F.; et al. ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education. *Learning and Individual Differences* **2023**, *103*, 102274.
9. Brynjolfsson, E.; McAfee, A. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*; W. W. Norton & Company, New York, 2014.
10. Frey, C.B.; Osborne, M.A. The Future of Employment: How Susceptible Are Jobs to Computerisation? *Technological Forecasting and Social Change* **2017**, *114*, 254–280.
11. Piaget, J. *The Origins of Intelligence in Children*; International Universities Press, New York, 1952.
12. Vygotsky, L.S. *Mind in Society: The Development of Higher Psychological Processes*; Harvard University Press, Cambridge, MA, 1978.
13. Lyman, F. The Responsive Classroom Discussion: The Inclusion of All Students. In *Mainstreaming Digest*; Anderson, A.S., Ed.; University of Maryland, College Park, 1981; pp. 109–113.
14. Flavell, J.H. Metacognition and Cognitive Monitoring: A New Area of Cognitive-Developmental Inquiry. *American Psychologist* **1979**, *34*, 906–911.
15. Hofer, B.K. Personal Epistemology as a Psychological and Educational Construct: An Introduction. *Personal Epistemology: The Psychology of Beliefs about Knowledge and Knowing* **2002**, pp. 3–14.
16. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; et al. Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys* **2023**, *55*, 1–38.

17. Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; et al. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *arXiv preprint arXiv:2311.05232* **2023**.
18. Wang, M.C.; Haertel, G.D.; Walberg, H.J. What Influences Learning? A Content Analysis of Review Literature. *The Journal of Educational Research* **1990**, *84*, 30–43.
19. Dignath, C.; Büttner, G.; Langfeldt, H.P. How Can Primary School Students Learn Self-Regulated Learning Strategies Most Effectively? A Meta-Analysis on Self-Regulation Training Programmes. *Educational Research Review* **2008**, *3*, 101–129.
20. Becher, T.; Trowler, P.R. *Academic Tribes and Territories: Intellectual Enquiry and the Culture of Disciplines*, 2nd ed.; Open University Press, Buckingham, 2001.
21. Colton, S.; Wiggins, G.A. Computational Creativity: The Final Frontier? *Frontiers in Artificial Intelligence and Applications* **2012**, *242*, 21–26.
22. Boden, M.A. *The Creative Mind: Myths and Mechanisms*, 2nd ed.; Routledge, London, 2004.
23. Popper, K.R. *The Logic of Scientific Discovery*; Hutchinson, London, 1959.
24. Mitchell, M. *Artificial Intelligence: A Guide for Thinking Humans*; Farrar, Straus and Giroux, New York, 2019.
25. Altieri, M.A. *Agroecology: The Science of Sustainable Agriculture*, 2nd ed.; CRC Press, Boca Raton, FL, 2018.
26. Bates, D.W.; Levine, D.M.; Salmasian, H.; Syrowatka, A.; Shahian, D.M.; Lipsitz, S.; et al. The Safety of Inpatient Health Care. *New England Journal of Medicine* **2023**, *388*, 142–153.
27. Charon, R. Narrative Medicine: A Model for Empathy, Reflection, Profession, and Trust. *JAMA* **2001**, *286*, 1897–1902.
28. Marmot, M.; Wilkinson, R.G. *Social Determinants of Health*, 2nd ed.; Oxford University Press, Oxford, 2005.
29. Fleddermann, C.B. *Engineering Ethics*, 4th ed.; Pearson, Upper Saddle River, NJ, 2012.
30. International Council for Harmonisation. Integrated Addendum to ICH E6(R1): Guideline for Good Clinical Practice E6(R2). ICH Harmonised Guideline, 2016.
31. Smith, C.M. Origin and Uses of *Primum Non Nocere*—Above All, Do No Harm! *The Journal of Clinical Pharmacology* **2005**, *45*, 371–377.
32. Vamathevan, J.; Clark, D.; Czodrowski, P.; Dunham, I.; Ferran, E.; Lee, G.; et al. Applications of Machine Learning in Drug Discovery and Development. *Nature Reviews Drug Discovery* **2019**, *18*, 463–477.
33. Cuban, L. *Oversold and Underused: Computers in the Classroom*; Harvard University Press, Cambridge, MA, 2001.
34. Paul, R.; Elder, L. *Critical Thinking: Tools for Taking Charge of Your Professional and Personal Life*, 2nd ed.; Pearson, Upper Saddle River, NJ, 2019.
35. Deakin Crick, R. *Learning to Learn: Setting the Agenda for Schools in the 21st Century*; Curriculum Journal, Taylor & Francis, 2008.
36. Cotton, D.R.E.; Cotton, P.A.; Shipway, J.R. Chatting and Cheating: Ensuring Academic Integrity in the Era of ChatGPT. *Innovations in Education and Teaching International* **2024**, *61*, 228–239.
37. Dewey, J. *Experience and Education*; Kappa Delta Pi, New York, 1938.