

THE TEACHER STILL DECIDES



A Practical Guide to AI
for the Higher Education Teacher

Adit Gupta

2026

OPEN ACCESS

*The Teacher Still Decides:
A Practical Guide to AI for the Higher Education
Teacher*

First published 2026
Copyright © Adit Gupta 2026

This work is licensed under a Creative Commons Attribution-Non-Commercial 4.0 International Licence (CC BY-NC 4.0).

You are free to:

Share — copy and redistribute this material in any medium or format.

Adapt — remix, transform, and build upon this material.

Under the following terms:

Attribution — You must give appropriate credit, provide a link to the licence, and indicate if changes were made.

Non-commercial — You may not use this material for commercial purposes without prior written permission from the author.

Full licence terms available at: <https://creativecommons.org/licenses/by-nc/4.0>



This publication is freely available for download at: <https://oar.miercollege.in>

Published by MIER College of Education (Autonomous), Jammu, India
ISBN: 978-81-952808-4-1

*The views expressed in this guide are those of the author and do not necessarily represent the views of any institution, organisation, or body with which the author is affiliated. Every effort has been made to ensure the accuracy of the information presented in this guide. The AI Interaction Knowledge (AIK) framework and the two-question test described in this guide were developed through research and practice at MIER College of Education, Jammu. The Continuous and Comprehensive Evaluation (C²E) model is an original institutional framework developed and implemented at MIER College of Education. This guide is a companion to *AI in the Classroom Before the Classroom: A Practitioner's Guide for Teacher Educators and Pre-Service Teachers in India* (Gupta, 2026), available from the same repository.*

Suggested citation:

Gupta, A. (2026). The teacher still decides: A practical guide to AI for the higher education teacher. MIER College of Education. <https://oar.miercollege.in>

This work is licensed under CC BY-NC 4.0. Full terms at creativecommons.org/licenses/by-nc/4.0

For permissions, institutional adoption, or correspondence:

adit@mier.in

Contents

Table of Contents

CHAPTER 1: BEFORE YOU BEGIN	1
THE QUESTIONS IN THE ROOM	1
WHERE DO YOU STAND? A SELF-ASSESSMENT	1
TWO MYTHS TO CLEAR AWAY	3
THE THINKING BEHIND THIS BOOK.....	4
THE FRAMEWORK: THE TWO-QUESTION TEST	4
WORKED EXAMPLE: THE TEST IN ACTION.....	5
 CHAPTER 2: PLANNING AND PREPARATION.....	 7
THE SUNDAY EVENING PROBLEM	7
THE FRAMEWORK: THE BRIEF, DRAFT, SHAPE, OWN CYCLE	7
WHAT PLANNING TASKS RESPOND BEST	9
WORKED EXAMPLE: THE CYCLE IN ACTION	9
 CHAPTER 3: CREATING TEACHING MATERIALS	 11
THE PHOTOCOPIER TEST	11
WHAT AI CREATES WELL.....	11
WHERE AI STUMBLES	11
THE FRAMEWORK: THE THREE FILTERS.....	12
WORKED EXAMPLE: THE FILTERS IN ACTION.....	13
 CHAPTER 4: TEACHING THE LARGE CLASS	 15
NINETY STUDENTS, ONE TEACHER.....	15
THE FRAMEWORK: ONE CONCEPT, THREE VERSIONS.....	15
THREE MORE LARGE-CLASS USES.....	16
WORKED EXAMPLE: THREE VERSIONS IN ACTION	17
 CHAPTER 5: ASSESSMENT IN THE AGE OF AI.....	 18
THE ESSAY THAT DIED.....	18

THE FRAMEWORK: PROCESS, LOCAL, LIVE.....	18
PROOF THAT IT SCALES: THE C ² E MODEL	20
WORKED EXAMPLE: ONE ASSIGNMENT, REDESIGNED	20
 CHAPTER 6: FEEDBACK AT SCALE.....	22
THE PILE ON THE DESK	22
THE FRAMEWORK: THE FEEDBACK BOUNDARY	22
THREE WAYS TO USE THE HONEST SIDE	23
WORKED EXAMPLE: NINETY ESSAYS, ONE BOUNDARY	24
 CHAPTER 7: YOUR STUDENTS ARE ALREADY USING IT	26
THE CONVERSATION NOBODY IS HAVING	26
THE FRAMEWORK: THE AI TRAFFIC LIGHT	26
THE CONVERSATION ITSELF	27
WORKED EXAMPLE: THE FIRST DAY, FIFTEEN MINUTES.....	28
 CHAPTER 8: YOUR OWN GROWTH.....	30
THE TEACHER WHO STOPPED LEARNING	30
THE FRAMEWORK: THE THREE ROLES	30
WORKED EXAMPLE: THE HISTORY TEACHER AND THE NUMBERS.....	31
 CHAPTER 9: AI IN YOUR RESEARCH AND PUBLICATION	33
WHERE THE STAKES CHANGE	33
THE ONE THING YOU MUST UNDERSTAND: INVENTED REFERENCES	33
THE FRAMEWORK: THE THREE ZONES	34
WORKED EXAMPLE: ONE PAPER, DONE HONESTLY	35
 CHAPTER 10: DOING THIS RESPONSIBLY	37
THE CHAPTER THAT KEEPS THE REST HONEST	37
THE THREE QUIET RISKS.....	37
THE FRAMEWORK: YOUR PERSONAL AI POLICY	38
WORKED EXAMPLE: ONE TEACHER'S PAGE	39
WHERE DO YOU STAND NOW?.....	39
 A CLOSING NOTE: THE DECISION THAT CANNOT BE DELEGATED	43

ABOUT THE AUTHOR	44
APPENDIX A: MY AI PRACTICE PLAN	45
APPENDIX B: THE PROMPT BANK.....	47
APPENDIX C: QUICK DECISION CARD	49

Chapter 1: Before You Begin

The questions in the room

Ask a room full of higher education teachers what they want to know about artificial intelligence, and the hands go up quickly. The questions, however, are rarely technical. In faculty development sessions across India, the same four surface every time.

Is using AI for our work cheating, or does it make us more productive?

Is using AI in research acceptable, or does it deviate from the standard research process?

How do we design assignments and assessments in the times of AI?

Does AI-generated content raise plagiarism concerns?

Notice what these questions share. Not one of them asks how to write a prompt or which tool to download. Every one of them is a question about legitimacy, integrity and professional identity. Underneath them sits a quieter question that few teachers voice aloud: *if a machine can do parts of my job, what happens to me?*

This book exists to answer that quieter question first, because until it is answered, no practical advice will stick. So let us answer it directly. Nothing in this book will ask you to hand over your professional judgement to a machine. Everything in this book rests on a single principle, and it is worth stating on page one: **the teacher still decides.**

Where do you stand? A self-assessment

Before reading further, take three minutes to answer honestly. For each statement, circle or tick one score: **Not yet = 1, Occasionally = 2, Regularly = 3, Confidently = 4.** Write your score for each item in the final column, then total it.

#	Statement	Not yet (1)	Occasionally (2)	Regularly (3)	Confidently (4)	My score
1	I can describe in one sentence what tools like ChatGPT, Gemini or Claude actually do.	1	2	3	4	

2	I can write a clear instruction to an AI tool and improve it when the first result disappoints.	1	2	3	4	
3	I check AI-generated content for errors before using any part of it.	1	2	3	4	
4	I can spot when an AI answer sounds confident but is actually wrong or invented.	1	2	3	4	
5	I have used AI to help prepare teaching materials such as questions, examples or outlines.	1	2	3	4	
6	I have thought about how AI changes the way I should design assessments.	1	2	3	4	
7	I have spoken with my students about acceptable and unacceptable uses of AI.	1	2	3	4	
8	I know which uses of AI in research are legitimate and which cross ethical lines.	1	2	3	4	
9	I am careful about what personal or student information I enter into AI tools.	1	2	3	4	
10	I can explain and defend any AI-assisted work I produce, with my name on it.	1	2	3	4	
	Total score					/40

Interpreting your score. 10 to 17, Beginning: This book was written precisely for you; read the chapters in order. **18 to 26, Experimenting:** The chapters on assessment, students and research will stretch you most. **27 to 34, Practising:** Use the book to make your habits deliberate rather than accidental. **35 to 40, Fluent:** Read for the frameworks and consider mentoring a colleague through the actions.

Keep your score safe. You will meet these ten statements again on the last pages of this book, and the distance between the two scores will be yours to measure.

Two myths to clear away

Almost everything unhelpful that is said about AI in education comes from one of two myths. Both are common, both sound plausible, and both are wrong.

The replacement myth says that AI will replace teachers. You have heard this in headlines, at conferences and possibly from worried colleagues. Here is why it fails. Teaching in higher education is not the delivery of information; if it were, the library would have replaced us long ago, followed by the internet. Teaching is the exercise of judgement: deciding what this group of students needs next, recognising the confusion behind a wrong answer, knowing when to push and when to encourage, and vouching for what counts as competence in your discipline. AI has no stake in your students' futures and no accountability for their learning. You do. A tool that cannot be responsible cannot replace someone whose entire role is responsibility.

The magic tool myth is the replacement myth's optimistic twin. It says that AI will do everything for you: write your lectures, mark your scripts, produce your research papers, all at the press of a button. This myth is more dangerous than the first, because it flatters rather than frightens. The truth is that AI output, left unchecked, is frequently generic, sometimes wrong and occasionally invented outright. A lesson plan generated in ten seconds is a starting point, not a finished product. An AI tool that produces a fluent paragraph with a fabricated reference has not saved you work; it has created work and risk that only your expertise can catch. Anyone who tells you AI removes the need for effort is selling something.

Reject both myths and a third position appears, which is where this book lives. AI is neither your replacement nor your servant. It is closer to a capable but inexperienced assistant: fast, tireless, occasionally brilliant, frequently overconfident, and entirely dependent on your supervision. Used this way, it genuinely does make you more productive, and no, that is not cheating. A teacher who uses AI to draft three versions of an explanation and then applies twenty years of classroom wisdom to choose and refine the best one is not cutting corners. That teacher is doing what professionals have always done with new tools: keeping the judgement and delegating the labour.

The thinking behind this book

There is a simple idea underneath every chapter that follows. Working well with AI is not mainly a technical skill; it is a form of professional knowledge, in the same way that knowing how to teach your subject is professional knowledge. In my research, I call this AI Interaction Knowledge, or AIK: knowing how to communicate with AI tools, how to judge and verify what they produce, how to fit them into your actual work, and how to use them responsibly. The ten statements you have just scored yourself against are, in fact, these four abilities in everyday language. You do not need to remember the term. You only need to notice that these are learnable skills, not talents, and that every chapter in this book builds one piece of them.

The framework: the two-question test

Each chapter of this book gives you one small framework to think with. Here is the first, and it answers the cheating question directly. Before using AI for any professional task, ask yourself two questions.

Question one: Am I still doing the thinking? If AI is generating options, drafts or summaries while you evaluate, select, correct and improve, you are working. If AI is producing finished work that you pass on without engagement, you are outsourcing, and outsourcing your judgement is where legitimate use ends.

Question two: Would I stand behind this output with my name on it? If you have checked it, shaped it and can defend every claim in it, the work is yours, whatever tools helped along the way. If you cannot explain or justify what you are submitting, presenting or publishing, then it is not yours, and no disclaimer will make it so.

Two yes answers mean you are on solid ground. Either no means stop and reclaim the task. This test will reappear throughout the book, in planning, in assessment and most sharply in Chapter 9, where we deal with research and publication. It works because it locates integrity where it has always lived: not in which tools you touch, but in whether the judgement and the accountability remain with you. Figure 1 presents the test as a simple decision path you can follow every time.




Figure 1: The Two-Question Test

Worked example: the test in action

The task: Dr Sharma, who teaches commerce, needs a notice informing students about a revised internal assessment schedule.

Her prompt: "Write a short, polite notice for undergraduate commerce students informing them that the internal assessment scheduled for 20 March is now postponed to 27 March because of the university sports meet. Mention that the syllabus coverage remains unchanged. Formal but warm tone, under 100 words."

The AI's output: a competent notice, correctly structured, but it invented a line stating that "the revised date-sheet is available on the university website", which was not true, and it addressed students as "Dear Learners", which is not how her institution writes.

Her judgement: she deleted the invented line, changed the salutation to her institution's convention, tightened one sentence and sent it. Total time: four minutes instead of fifteen.

The test applied: Was she still doing the thinking? Yes, she specified the content, caught the fabrication and enforced house style. Would she stand behind it with her name on it? Yes, every line was verified. Legitimate use, genuine productivity, and a small lesson learnt: AI invents details confidently, even in a hundred-word notice. Imagine what it can invent in a research paper. We will return to that in Chapter 9.

Try this week

Take one routine task from your week, perhaps a notice, a quiz question or a summary of a reading. Ask any freely available AI tool to attempt it, using a prompt that states the audience, purpose, tone and length, as Dr Sharma did. Then examine the output rather than adopting it. Find at least one thing it got wrong or missed about your context; there may be none. Fix it, apply the two-question test, and only then use it. Fifteen minutes, no cost, no permission needed, and you will have learnt more about what AI is and what it is not than any headline could teach you.

The takeaway

AI can produce the words; you provide the judgement. Every chapter that follows simply applies this one principle to a different part of your work. The teacher still decides.

Chapter 2: Planning and Preparation

The Sunday evening problem

Every teacher knows Sunday evening. The new week needs a lecture outline, a tutorial task, discussion questions for a seminar and perhaps a revised unit plan, and all of it competes with the marking you carried home and the committee report due on Tuesday. Planning is the most invisible part of teaching: students never see it, appraisals rarely measure it, and yet it quietly consumes more hours than the classroom itself. It is also, not coincidentally, the safest and most rewarding place to begin using AI, because everything AI produces at the planning stage passes through your hands before any student ever sees it. Here, there is no risk of a machine addressing your students directly. There is only you, a drafting assistant of unusual speed, and your own judgement standing between the two.

But there is a trap waiting, and it is worth naming before the first prompt is typed. Ask an AI tool for "a lesson plan on marketing basics" and you will receive one in seconds: neat, plausible and utterly generic. It will assume students it has never met, a curriculum it has never read and a classroom it cannot imagine. Teachers who stop at this point conclude that AI produces bland, unusable material, and they are half right. The output was bland because the input was empty. AI does not know that your second-year class has forty-five students of wildly mixed English proficiency, that your university follows a prescribed syllabus, that your teaching block is fifty minutes, or that Tuesday's class falls immediately after lunch when energy is lowest. You know all of this. Planning well with AI is therefore not about clever tricks; it is about transferring your context into the conversation before asking for anything.

The framework: the Brief, Draft, Shape, Own cycle

Every planning task with AI follows the same four steps, and noticing them turns a vague chat into a professional method.

Brief: You begin, not the machine. Tell it who the learners are, what the syllabus requires, how much time you have, what resources exist and what problem you are trying to solve. A good brief is five or six sentences of pure context. This step is entirely yours, and it is where most of the quality is determined.

Draft: The AI produces material: an outline, activities, questions, and a reading structure. Ask for more than one version; three variations cost nothing and give you choices instead of a verdict. This step belongs to the machine, working within the walls your brief built.

Shape: You now cut, reorder, correct and localise. Replace the generic example with one from your city; remove the activity that needs technology your classroom lacks; fix the factual slip; adjust the timing you know is unrealistic. Shaping usually takes longer than briefing, and it should, because this is where the plan becomes yours.

Own: Apply the two-question test from Chapter 1. Are you still doing the thinking? Would you stand behind this plan with your name on it? If yes, the plan goes into your file as your work, because it is. If not, return to shaping.

The cycle matters because it makes visible where the teacher sits: at the beginning, deciding what is needed, and at the end, deciding what is good. The AI tool occupies only the middle. Figure 2 shows the four steps and who owns each of them; skip the first step, and you get generic output; skip the last two, and you are no longer planning, merely accepting.




Figure 2: The Brief, Draft, Shape, Own Cycle

What planning tasks respond best

Three preparation tasks give the fastest returns for a higher education teacher. The first is **structuring**: turning a syllabus topic into a sequenced outline with timings or converting fourteen weeks of prescribed content into a coherent teaching calendar. AI is genuinely strong at structure, and structural output is easy to check because you can see at a glance whether the sequence makes sense. The second is **generating variety**: five discussion questions at different levels of difficulty, three alternative openings for the same lecture, four real-world applications of one concept. Variety is expensive for a tired human brain and nearly free for a machine, and you remain the one who selects. The third is **anticipating**: ask what misconceptions students typically hold about a topic, or what questions a confused student might ask, and you receive a rehearsal partner for the classroom. None of these tasks lets the AI near your students; all of them return an hour of your Sunday evening.

One planning task deserves a caution flag: **reading lists**. AI tools can and do invent books, articles and authors that sound entirely real and do not exist. Never accept a suggested reading you have not personally verified in your library catalogue or a database. Use AI to suggest the shape of a reading list, themes and types of sources, but confirm every individual item yourself. This is the first appearance of a warning that will return with full force in Chapter 9.

Worked example: the cycle in action

The task: Dr Verma teaches an undergraduate psychology paper and needs a 50-minute introductory session on memory for a first-year class of about 60 students, with mixed English proficiency; no projector available that week.

Her brief: *"I teach first-year undergraduate psychology in India. I need a plan for one 50-minute introductory class on human memory for about 60 students with mixed English proficiency. No projector or slides available, only a blackboard. The prescribed textbook covers sensory, short-term and long-term memory. I want at least one activity that gets students participating without needing any materials. Give me two different plan options with rough timings."*

The drafts: the AI returned two structured plans with timings. Option one opened with a whole-class memory game requiring no materials; option two opened with a lecture segment. Both were usable; both had flaws. One plan allotted ten minutes for an activity that would clearly take twenty with sixty students, and one suggested "showing a short video", ignoring her stated constraint.

Her shaping: she chose the game-based opening from option one, doubled its time allocation, deleted the video suggestion, replaced an American example about eyewitness

testimony with a locally relevant one, and merged the closing recap from option two, which was stronger.

Her ownership: the final plan reflected her class size, her room, her textbook and her instinct for what sixty first-years can do after lunch. Total planning time: about twenty minutes, including the shaping. Was she still doing the thinking? Entirely. Would she put her name on it? She taught it on Tuesday.

Try this week

Choose one class you must teach next week and write a brief for it before asking anything of an AI tool: learners, class size, time, syllabus point, room constraints and one problem you want solved. Make the brief at least five sentences. Then ask for two plan options and deliberately find three things to change during shaping. The point of the exercise is not the plan; it is discovering how directly the quality of the brief controls the quality of the draft. Teachers who learn briefing in week one never write "give me a lesson plan on marketing basics" again.

The takeaway

The AI tool can draft a plan in seconds, but only you know the classroom it must survive in. Brief with your knowledge, shape with your judgement, and the resulting plan is yours. The teacher still decides.

Chapter 3: Creating Teaching Materials

The photocopier test

Walk past the photocopier in any college, and you will see the real curriculum: handouts, question banks, case studies, tutorial sheets. These materials teach quietly long after your lecture ends. They also consume your evenings.

AI can draft all of them quickly, and like planning, this work happens before any student is involved. But materials carry one special risk. A weak lecture harms one class and is forgotten. A handout with an error is photocopied sixty times and revised from before the examination. Materials multiply, and their mistakes multiply with them. So, the standard here must be higher than anywhere else in your preparation.

What AI creates well

Four kinds of material give the best returns.

Handouts and summaries: Give the AI your lecture outline and ask for a one-page summary at a stated reading level.

Question banks: Multiple-choice, short-answer, and application questions, in quantity and at varying difficulty levels.

Case studies: Short, realistic situations that turn a concept into a decision students can argue about.

Worked examples: Step-by-step solutions with the reasoning shown. You can even ask for a deliberately flawed solution for students to critique.

Where AI stumbles

Three weaknesses appear again and again. Know them in advance and they lose most of their power.

Foreign defaults: Case studies set in California, prices in dollars, examples your students have never seen.

Factual slips: Wrong facts, wrong calculations, and most dangerously, multiple-choice questions where the AI marks the wrong option as correct. An error in a lecture can be corrected on the spot. An error in an answer key silently mismarks the whole class.

Difficult language: The default vocabulary often runs well above the comfortable reading level of students who work in English as a second or third language.

None of this makes AI drafting a bad idea. It simply gives you your checklist.

The framework: the three filters

Chapter 2 gave you the Brief, Draft, Shape, Own cycle. When the product is teaching material, the three filters are what Shape means. Every draft passes through all three before it reaches a student. Figure 3 shows them as a funnel between the AI's draft and your classroom.

Filter 1, Accuracy: Is every fact, figure and answer correct? Check each claim. Recalculate each example. Attempt every quiz question yourself before trusting its answer key. Never ask the AI to check its own work; the system that made the error cannot be trusted to find it.

Filter 2, Context: Does it fit my students and my place? Replace the foreign company with a local one, the dollar with the rupee, the unfamiliar scenario with one your students walk past daily.

Filter 3, Language: Can my weakest reader follow it? If not, ask for shorter sentences and simpler words. A handout only your strongest students can read is a handout for part of your class.

The order matters. Accuracy first, because wrong content ruins everything else. Context second, because relevance decides whether students engage. Language last, because polish means nothing on material that is wrong or foreign.




Figure 3: The Three Filters

Worked example: the filters in action

The task: Mr Qureshi teaches introductory economics to first-year B.Com students. He wants a short case study on demand and supply, with five multiple-choice questions.

His prompt: *"Write a one-page case study for first-year commerce students in India about demand and supply, based on a vegetable seller in a weekly market whose tomato prices change through the day. Simple English, short sentences, prices in rupees. Then create five multiple-choice questions on the case, with four options each, and indicate the correct answers."*

The draft: Usable and correctly set in a weekly market with rupee prices, as his brief demanded. The filters still caught three problems.

Accuracy: one question marked option B as correct when the case's own numbers pointed to option C. He recalculated and fixed the key. **Context:** the seller was "checking prices on a commodity exchange app"; he changed it to a phone call to the mandi. **Language:** one paragraph read "equilibrium is restored through price signalling mechanisms"; he asked for a rewrite in basic English and the second version was clean.

The result: a local case study and a verified question bank in under thirty minutes, most of it spent filtering, exactly as it should be.

Try this week

Ask an AI tool for five multiple-choice questions, with answers, on a topic you will teach this fortnight. Attempt all five yourself before looking at its key. If every key is right, you have seen what a well-briefed AI can do. If even one is wrong, you have learnt, at no cost, why Filter 1 exists. Keep the good questions; your question bank has started.

The takeaway

AI can draft an afternoon's worth of materials in minutes, but the photocopier multiplies whatever you feed it, errors included. Filter for accuracy, context and language, and what reaches your students carries your standard, not the machine's. The teacher still decides.

Chapter 4: Teaching the Large Class

Ninety students, one teacher

The timetable says one class, but you know better. In a room of ninety, there are really three classes sitting together: students who will follow your first explanation, students who need a second one in simpler words, and students who are quietly lost by minute ten and too shy to say so. Add mixed English proficiency and mixed school backgrounds, and the large Indian classroom becomes the hardest differentiation problem in education.

The honest truth is that AI cannot fix class size. It will not reduce ninety to thirty, and it will not learn your students' faces. What it can fix is the arithmetic of preparation. Differentiation has always failed in large classes for one reason: preparing three versions of everything triples the workload, and no teacher has three evenings for every lecture. AI collapses that cost. Preparing for one class or preparing for the three classes hiding inside it now takes nearly the same time. That changes what is possible.

The framework: one concept, three versions

For every difficult concept in your course, prepare three versions before you walk in. Figure 4 shows the idea: one concept, three doors into it.

The simple version: Plain language, short sentences, one everyday analogy. This is for students meeting the idea for the first time, and for students whose English makes the textbook a wall. It is not a lesser version; it is the entry point.

The standard version: The textbook-level explanation with a worked example, matching what the examination will expect. Most of the class lives here.

The stretch version: An application, a twist or a problem without an obvious answer, for students who finished early and are getting bored. Boredom in the front rows is as real a teaching failure as confusion in the back rows.

Ask the AI to draft all three at once; it is precisely the kind of task it does well. Your job is to pass each version through the three filters from Chapter 3 and to decide when each version enters the lecture. Typically, the simple version opens, the standard version carries the middle, and the stretch version becomes the closing question or the take-home challenge.




Figure 4: One Concept, Three Versions

Three more large-class uses

A bank of varied examples: In a class of ninety, one example never lands with everyone. Ask for five examples of the same concept drawn from different spheres: a family, a shop, a farm, a cricket match, a phone. When one example misses a student, the next one catches them.

Bilingual support: Ask the AI for a glossary of the topic's technical terms with plain-English meanings and equivalents in your students' stronger language. For many students, the concept is not the barrier; the vocabulary is. A half-page glossary, checked by you, removes that barrier for the cost of one prompt.

Quick checks that fit ninety: You cannot ask ninety students individual questions, but you can end class with one well-designed question everyone answers on a slip of paper. Ask the AI to draft one-minute-paper questions or a five-item true-and-false check for today's topic. Reading ninety slips takes ten minutes and tells you more than silence ever will.

Worked example: three versions in action

The task: Dr Bhat teaches statistics to eighty-five B.A. psychology students with very mixed mathematics backgrounds. Next week: standard deviation, the lecture where she loses people every year.

Her prompt: *"I teach statistics to first-year psychology students in India, many of whom fear mathematics. Explain standard deviation in three versions: first, a simple version using an everyday analogy and no formula; second, a standard textbook version with one small worked example using marks of five students; third, a stretch version posing a thought question that applies the idea to psychology. Simple English throughout."*

The drafts: the simple version compared two cricket batsmen with the same average but very different consistency, which she knew would land. The standard version's worked example was correct; she recalculated it anyway, filter one being non-negotiable. The stretch version asked why two anxiety questionnaires with the same mean score might describe very different groups of patients, which was exactly the bridge to psychology she wanted.

Her shaping: she swapped one batsman's name for a current player her students follow, and asked for a ten-term English-to-Hindi glossary of the lecture's statistical vocabulary as a handout.

The class: the cricket analogy opened, the worked example carried the middle, the questionnaire puzzle closed, and the glossary sat in ninety bags. Preparation time: forty minutes for a lecture that used to fail annually.

Try this week

Pick the one concept in your next fortnight that students historically find hardest; every teacher knows theirs instantly. Ask an AI tool for the three versions, using a prompt like Dr Bhat's, with your subject and your students in it. Filter all three, then teach the concept with the simple version first. Watch the back rows during the analogy. That is where you will see whether the third door opened.

The takeaway

AI cannot shrink a class of ninety, but it can end the era of one explanation for everyone. Prepare three versions at the cost of one, and every student finds a door into the concept. The teacher still decides.

Chapter 5: Assessment in the Age of AI

The essay that died

Set this assignment today: *"Discuss the impact of social media on Indian youth. 1,500 words, submit in two weeks."* A free AI tool completes it in forty seconds, at a quality that would have earned a comfortable B two years ago. Your students know this. The take-home essay, the workhorse of higher education assessment for a century, no longer measures what it used to measure. Pretending otherwise is the one option a serious teacher does not have.

Faced with this, institutions reach for two responses, and both fail.

Banning AI fails because a ban you cannot enforce is not a rule; it is a lottery that punishes only the honest.

Detection fails because no AI detector is reliable. Detectors miss real AI text, and worse, they flag genuine student writing as machine-made, most often the writing of students working in their second language. An accusation built on a faulty detector can damage a student for life. Do not build integrity on that foundation.

There is a third response: **redesign**. Assume every student has AI access, the way we long ago assumed every student has a calculator, and then design assessments that still measure learning under that assumption. Redesign asks more of the teacher than banning and is more honest than detecting. It is also the only response that works.

The framework: Process, Local, Live

Not every assessment needs changing; a supervised examination hall is already AI-proof. The vulnerable genre is unsupervised take-home work, and for it there are four redesign moves: three that restructure the task, and a bold fourth that welcomes AI into it openly. Figure 5 shows how these moves transform a vulnerable task. You rarely need all four at once; even one changes the game.

Move one: assess the Process, not only the product. Require the stages, not just the destination: a topic proposal, an outline, a draft with your comments answered, then the final piece. AI can produce a product in forty seconds; it cannot fake a month-long trail of developing thought.

Move two: anchor in the Local. Tie the task to material AI has never seen: this week's class discussion, data the student collects, an interview with a real person, the market outside

the campus gate. "Discuss the impact of social media on youth" is AI's home ground. "Interview three students of this college about their social media use and analyse your findings using the two theories from Unit 2" is yours.

Move three: add a Live element. Attach a short viva, a presentation or an in-class application task to the written submission. Five minutes of "explain your third paragraph" reveals instantly whether the mind behind the work is present. In a class of ninety, viva time is real, but even a random viva for one student in five changes behaviour for all ninety, because nobody knows who will be called.

Move four: Open the door. For some tasks, allow AI explicitly and assess what the student does with it. Permit its use, require evidence of how it was used, and mark the thinking the student added beyond the machine. This is the boldest move, and it deserves more than a paragraph, because I have watched it work at an institutional scale.




Figure 5: Redesigning the Vulnerable Task.

Proof that it scales: the C²E model

At my own institution, we stopped patching individual assignments and redesigned the assessment system itself. The result is the Continuous and Comprehensive Evaluation (C²E) Four-Assessment Model, under which every theory paper is assessed through four components. Each component is one of this chapter's principles turned into policy.

A1, the Analytical Assignment, allows AI openly. Students may use AI tools, but they must submit their chat history along with the assignment as evidence of how they used them. What earns marks is the thinking the student added beyond the machine, and the chat record makes that thinking visible. This is the two-question test from Chapter 1, taught to students by design rather than by lecture.

A2, the Open Book Examination, assumes access instead of policing it, but draws the line with precision. Students may bring books, notes and any paper-based material into the hall; what they may not bring is anything electronic or connected, no laptops, no mobiles, no wearables. The distinction is deliberate. Paper can remind a student of what a concept is; it cannot apply that concept to the unfamiliar scenario on the question paper. The examination therefore tests exactly what remains valuable when information is freely available: judgement under supervised conditions.

A3, the ICT-Based Quiz, is timed and supervised. It protects the one thing still worth protecting the old way: foundational knowledge held in the student's own head.

A4, the Skill Micro-credential, sends students to complete a verified online course linked to the paper, through platforms such as LinkedIn Learning. The certificate alone, however, is not enough: each student also records a two-minute reflective video explaining what they learnt and how it connects to their course. A certificate can be earned mechanically; a two-minute explanation in your own words cannot. The video is, in miniature, the same principle as the viva: the fastest way to verify learning is to ask the learner to speak about it.

Worked example: one assignment, redesigned

The old task: Dr Nair, teaching sociology, has set the same term paper for years: *"Discuss the changing nature of the Indian family. 2,000 words."* Last term, a third of the submissions read suspiciously alike, and she could prove nothing.

The redesign: she kept the topic and applied the three moves. The new task: *"Interview one person above sixty and one person below twenty-five in your own family or neighbourhood about how family life has changed. Submit your interview notes (week 4), a one-page analysis plan (week 7), and a final 1,500-word paper connecting your interviews to the two theories studied in class (week 12). Five students each week will present their findings for five minutes."*

What changed: AI can still polish the final paper's language, and Dr Nair does not mind; language polish was never what she was marking. But no AI has met the student's grandmother. The interviews are Local, the staged submissions are Process, the presentations are Live, and the task is now more educational than the original, not merely more secure.

The honest cost: redesign took her one evening, with an AI tool drafting the staged structure and the rubric, which she then filtered. One evening, for an assignment that measures again.

Try this week

Take your most AI-vulnerable assessment; you already know which one it is. Ask an AI tool to complete it, exactly as a student would. Read the result once, honestly. If it passes, the assessment is measuring nothing. Redesign it with one move, whichever fits your subject, and let AI itself draft the new version and rubric for you to filter. The tool that broke your assignment can help you rebuild it, and there is a certain justice in that.

The takeaway

Bans cannot be enforced and detectors cannot be trusted, so the only real answer is design. Assess the process, anchor in the local, add a live element, and where you allow AI, make it visible; that is how assessment measures learning again. The teacher still decides.

Chapter 6: Feedback at Scale

The pile on the desk

Every teacher believes in feedback; the pile of ninety scripts on the desk believes in nothing. Research has said for decades that specific, timely feedback is among the most powerful influences on learning, and every teacher knows it from experience. But arithmetic is merciless. Five minutes of written comments per script, times ninety, is seven and a half hours, on top of the marking itself. So feedback quietly shrinks to what scale allows: a tick, an underline, a grade, "good work" or "needs improvement". Students learn their marks and nothing else. The problem was never the teacher's commitment; it was always the arithmetic. AI changes the arithmetic, but only if one boundary is respected absolutely, and this chapter is really about that boundary.

The framework: the feedback boundary

Some parts of feedback are judgement, and some parts are communication. The boundary between them determines what may be shared with a machine, and Figure 6 illustrates it clearly.

Always yours, never delegated: reading the student's work, judging its quality, and deciding the mark. These three are the professional act itself. An AI has not sat in your classes, does not know what you taught, and carries no accountability for a mark that shapes a student's future. The moment a machine assigns marks, assessment has ended, and something else has begun.

Where AI can honestly help: phrasing your comments clearly and kindly, compiling the errors you found into a whole-class feedback sheet, and building a reusable bank of well-written comments. These are communication tasks. The judgement in them is already yours; the machine only helps it travel. The whole boundary fits in one line: let AI help you say it, never decide it.

One rule guards the boundary, and it is absolute: **no student's name, roll number or identifiable work goes into an AI tool.** Student work belongs to the student. Describe the error, not the person: "a student argued X without evidence" carries everything the AI needs and nothing it should not have. Chapter 10 returns to why this matters; for now, treat it as a professional reflex.




Figure 6: The Feedback Boundary

Three ways to use the honest side

Whole-class feedback. This is the single best trade of effort for impact at scale. While marking, keep a rough tally of the five most common errors in your own shorthand. Then brief the AI: here are the five errors my class made on this assignment; draft a one-page feedback sheet that explains each error, shows a corrected example, and keeps an encouraging tone. Filter it, print it, and every one of ninety students receives substantial feedback for the cost of one prompt and one photocopy run.

A comment bank in your own voice. You have written "your argument needs evidence" four hundred times in your career. Give the AI your 10 most repeated comments and ask for 3 variations of each: one gentle, one direct, and one for a strong student who can take a push. Review them once and save them in a document. During marking, you simply pick the right

comment for each student and adjust a word or two. The words come from the bank, but the choice of which student gets which comment is still your judgement at work.

Feed-forward for drafts. Teach students to ask AI for feedback on their drafts before submission, against the assignment's actual rubric, which you share with them. The AI critiques structure and clarity; the student revises; what reaches your desk is better, and your feedback can then address thinking rather than tidiness. This costs you nothing and quietly teaches students the legitimate use of AI, which Chapter 7 takes up directly.

Worked example: ninety essays, one boundary

The task: Dr Joshi has eighty-eight first-year literature essays on her desk and a two-week deadline she has never once met.

Her process: she read and marked every essay herself, as always, but kept a running tally as she went. Four errors dominated: plot summary in place of analysis, quotations dropped in without comment, conclusions that introduced new points, and paragraph-long sentences.

Her prompt: *"I teach first-year English literature in India. My class of 88 made these four errors in their essays on 'The Postmaster': [her four notes, in her own words]. Draft a one-page whole-class feedback sheet. For each error, name it plainly, explain why it weakens an essay, and show a two-line corrected example using a generic sentence rather than the story. Warm, encouraging tone. Simple English."*

The boundary held: no essay, no name, no student writing entered the tool, only her own four observations. The marks were decided before she typed a word of the prompt.

The result: essays returned in nine days instead of the usual four weeks, each with its mark and brief margin notes, plus the feedback sheet discussed for fifteen minutes in class. Several students later said it was the first time they understood what "analyse, don't summarise" actually meant. The feedback got better and faster at the same time, which the old arithmetic never allowed.

Try this week

Next time you mark any set of work, keep a scrap of paper beside you and tally the repeated errors; it costs nothing while you mark. Then spend ten minutes turning your tally into a whole-class feedback sheet with an AI tool, filtered through Chapter 3's three filters before it is photocopied. One sheet, one class discussion, ninety students better informed than a term of ticks ever made them.

The takeaway

Reading, judging, and marking are yours and must stay yours; phrasing, compiling, and formatting are labour the machine can carry. Let AI help you say it, never decide it, and feedback finally scales without losing its soul. The teacher still decides.

Chapter 7: Your Students Are Already Using It

The conversation nobody is having

Ask your students privately whether they use AI for their coursework, and most will say yes. Ask them whether any teacher has ever discussed with them how to use it well, and most will say no. That gap is the real problem. In most classrooms, AI use has become something students do quietly and teachers pretend not to see. Nobody is learning anything from this silence. Students are not learning judgment because nobody is teaching it. Teachers are not learning what their students actually do, because nobody dares to ask.

Silence also creates unfairness. Your confident students are already using AI skilfully to explain difficult topics, to check their drafts, and to prepare for examinations. Your struggling students are either misusing it, submitting work they do not understand, or avoiding it entirely out of fear. The gap between the two groups grows every month that the classroom stays silent. The single most useful thing a teacher can do about student AI use costs nothing: talk about it, openly and early.

The framework: the AI traffic light

Students misuse AI most often for one simple reason: no one has told them what is allowed. "Do not cheat" is not guidance, because a student genuinely cannot tell whether asking AI to explain a concept, improve a paragraph, or write a conclusion counts as cheating. The rules live in the teacher's head, and each teacher's rules are different.

The traffic light replaces this confusion with three clear colours. Every task you set gets one colour, declared in advance, and Figure 7 shows the whole system on a single page.

Red-AI not permitted: Reserved for tasks whose entire purpose is to build a skill the student must carry in their own head, such as foundational writing exercises, mental calculation, or preparation for supervised examinations. When you set a red task, say why: "This task is red because the skill it builds is the one you will need in the examination hall, where you will be alone."

Amber-AI permitted with disclosure: The student may use AI, but must say how, in a short note attached to the submission: "I used AI to explain the concept of elasticity and to check my grammar." Amber is where most take-home work belongs, and the disclosure note is where students practise honesty about their tools. If your institution follows a model like C²E from Chapter 5, the chat-history evidence of the Analytical Assignment is simply amber taken seriously.

Green-AI expected and encouraged: Some tasks should require AI, because learning to use it is part of the point: "Get AI to generate three summaries of this article, then write one paragraph on which summary is best and why." Green tasks teach the evaluation skills of Chapter 1's two-question test, with the student in the examiner's chair.

The colours do something subtle and valuable: they replace suspicion with clarity. A student can no longer say "I did not know", and a teacher no longer has to guess.




Figure 7: The AI Traffic Light

The conversation itself

The traffic light needs fifteen minutes of honest talk to launch it, ideally in the first week of the course. Three things belong in that conversation.

Say the true thing first: "I know many of you use AI. I use it too. This course will not pretend otherwise." Nothing you say afterwards will land if you begin with a pretence that students know to be false.

Give them the gym analogy: A machine at the gym can lift the weight for you, and the weight will move, but your muscles will gain nothing. AI can write the essay for you, and the essay will exist, but your mind will gain nothing. The purpose of coursework was never the essay; it was always the strength you build by writing it. Students accept this reasoning readily because it treats them as people making choices rather than suspects being warned.

Explain the colours and the disclosure note: Show them the course outline with each task's colour marked. Tell them the disclosure note will never be punished; only its absence will. The moment one student is penalised for honestly declaring AI use, honesty in your classroom ends.

One more situation belongs in this chapter, because every teacher will face it: the submission you suspect was written by AI on a task where it was not permitted. Do not reach for a detector; Chapter 5 explained why they cannot be trusted. Reach for a conversation instead. Sit with the student and ask them to walk you through their work: how they chose the argument, what the third paragraph means, what they would change now. A student who did the work can talk about it, imperfectly but genuinely. A student who did not will show you within minutes, and the conversation itself, handled without humiliation, often teaches more than any penalty.

Worked example: the first day, fifteen minutes

The task: Dr Fernandes teaches a research methods course and decided to open the semester differently this year.

What she did: she spent fifteen minutes of the first lecture on AI. She began by telling the class that she uses AI herself to draft materials and check her own writing, and she showed them one example. Then she gave the gym analogy. Finally, she distributed the course outline, with every assessment task carrying a colour: the weekly reflections were red, with the reason stated; the literature review was amber, with a sample disclosure note printed on the outline itself; and one task was green, requiring students to critique an AI-generated research design.

What happened: disclosure notes arrived on nearly every amber submission, some charmingly detailed. Two students asked whether using AI to translate their thoughts from Hindi to English needed disclosure, a genuinely good question she had not considered, and the class discussed it together and decided it belonged in the note. The suspicious-submission conversations she used to dread did not happen that term. There was nothing left to hide.

Try this week

Choose the next assignment you plan to give and assign it a colour. Write the colour and one sentence of reasoning at the top of the assignment sheet. If it is amber, include a two-line sample disclosure note so students can see what honesty looks like. Then spend ten minutes of class explaining it. You will learn more about your students' actual AI use from that one conversation than from a year of wondering.

The takeaway

Students are already using AI; the only question is whether they use it in silence or with guidance. Give every task a colour, welcome the disclosure note, and replace suspicion with clarity. The teacher still decides.

Chapter 8: Your Own Growth

The teacher who stopped learning

There is a quiet irony in our profession. We spend our working lives helping other people learn, and our own learning gets whatever time is left over, which is usually none. The occasional faculty development programme arrives once or twice a year. Between those islands, growth depends on stolen evenings, and stolen evenings usually go to marking. Meanwhile, the questions pile up: the statistical method you never properly learnt, the teaching approach everyone mentions but nobody explains, the new development in your own discipline that you nod along to in seminars while privately hoping nobody asks you about it.

AI changes this situation in a way that has nothing to do with saving time on paperwork. For the first time in your career, you have access to a tutor who is available at midnight, never runs out of patience, never judges the question, and will explain anything at exactly the level you need. Every teacher has questions they have been too embarrassed to ask for twenty years. That embarrassment was the price of learning in front of colleagues. In front of a machine, the price drops to zero.

The framework: the three roles

AI can serve your growth in three distinct roles, and it helps to know which one you are asking it to play. Figure 8 shows all three side by side.

The Explainer: This is the role most teachers discover first. Ask it to explain any concept, at any level, in any style: "Explain factor analysis to someone who is good with words but afraid of mathematics." If the explanation is still too difficult, say so, and ask for a simpler one. If it is too simple, ask it to go deeper. You can ask the same question five times in five ways, which no human tutor's patience allows. This is the role where the judgement-free quality matters most: the question you have avoided asking for years takes two minutes to finally ask.

The Coach: In this role, the AI practises with you. Before teaching a new topic, explain it to the AI in your own words and ask it to find the gaps in your explanation. Before a conference presentation, ask it to play a difficult audience member and question you. Before a class discussion on a controversial topic, ask it to argue the side you find hardest to argue. One warning makes this role work: AI is agreeable by default, and if you simply ask "is my explanation good?", it will tell you yes. You must ask it directly to challenge you: "Point out the weaknesses in this explanation. Do not be polite." The critique that comes back is often genuinely useful, but only if you demand it.

The Scout: In this role, the AI surveys territory you have not visited: "What are the main approaches researchers currently use to study classroom interaction?" or "What is this new method my colleague mentioned?" The Scout is useful and dangerous in equal measure, because this is where AI most confidently presents outdated information, oversimplified summaries and, occasionally, things that are simply not true. Treat everything the Scout brings back as a rough map drawn from memory, helpful for orientation, but never to be trusted for the final route. Verify anything important against real sources before you build on it, a discipline that Chapter 9 turns into a firm rule.




Figure 8: The Three Roles

Worked example: the history teacher and the numbers

The situation: Mr Thomas has taught history for eighteen years and is respected for it. But his postgraduate students increasingly want to use surveys and simple statistics in their dissertations, and he has spent years quietly steering them away from quantitative work, not

because it was wrong for their topics, but because he could not supervise it. He has never admitted this to anyone.

The Explainer: he began by asking an AI tool the question he could never ask a colleague: "Explain what a chi-square test actually tells you, for someone who has avoided statistics for twenty years. No formulas yet." Twenty minutes of follow-up questions later, including several he described as embarrassingly basic, the concept was clearer than any workshop had ever made it.

The Coach: a week later, he tested himself. He explained chi-square back to the AI in his own words and asked it to find the errors, adding "do not be polite". It found two genuine misunderstandings, which he then fixed.

The Scout: finally he asked what quantitative methods historians actually use, received a promising overview, and then did the essential thing: he checked the suggested approaches against two methodology books in the library, and found one AI suggestion that did not match how historians really work. The map was useful; the verification was necessary.

The result: this term, for the first time, he approved a student's survey-based dissertation proposal, and he is supervising it with growing confidence. Nothing about his teaching load changed. What changed is that his own learning no longer waits for a workshop that never comes.

Try this week

Write down the one concept in your professional life that you have pretended to understand for years; every honest teacher has one. Spend twenty minutes with an AI tool asking about it, starting with "explain this simply" and following up with every question you have ever swallowed. Nobody is watching, nobody is judging, and twenty years of polite nodding can end in one evening.

The takeaway

Your students are not the only learners in your classroom, and your growth should not have to wait for the next development programme. Use the Explainer without embarrassment, push the Coach to criticise you, and verify everything the Scout brings home. The teacher still decides.

Chapter 9: AI in Your Research and Publication

Where the stakes change

Everything in this book so far has concerned your classroom, where mistakes are recoverable. A weak lesson plan dies quietly; a flawed handout can be corrected next week. This chapter concerns the one part of your work where mistakes are not recoverable. In Indian higher education, the teacher and the researcher are the same person: publications shape your appraisals, your promotions and your standing, and research supervision shapes your students' futures. A published paper carries your name permanently. There is no next week in which to correct it.

I write this chapter from a particular vantage point. As a journal editor, I now regularly see what AI misuse in research looks like when it arrives at the other end: manuscripts written entirely by machine and barely read by their own authors, reference lists citing papers that do not exist, and reviews that were plainly never written by the reviewer. Editors have learnt to recognise all of these, and the consequences fall on the author, not the machine. So, this chapter is firmer than the others, and it should be.

Let one thing be said clearly before the rules, though: AI has a legitimate and genuinely valuable place in research. The same tool that can end a career when misused can, when used honestly, help you search literature faster, sharpen your writing, and think through your analysis. The purpose of this chapter is not to frighten you away from the tool. It is to draw the map so precisely that you never wander into the wrong territory by accident.

The one thing you must understand: invented references

Before the framework, one technical fact explains most research disasters with AI, and every researcher must understand it.

When you ask an AI tool for references on a topic, it does not look anything up. It generates text that has the shape of references: plausible authors, a believable title, a real-sounding journal, a year, even a DOI. Sometimes these correspond to real papers. Often they do not. The tool is not lying in any deliberate sense; it is doing what it always does: producing plausible text. A fabricated reference is a plausible text of a particularly dangerous kind because it looks exactly like a real one. It cannot be spotted by reading. It can only be caught by checking.

This is why the rule that follows has no exceptions. **Every reference in your work must be verified by you: the paper found in a real database, opened, and actually read before it is cited.** A reference list is a set of promises to your reader that these sources exist and say

what you claim. An unverified reference is a promise you have no right to make. And to close one tempting loophole: asking the AI "are these references real?" is not verification. The system that invented them will cheerfully vouch for them.

The framework: the three zones

Every use of AI in research falls into one of three zones, and Figure 9 maps all three. The zones are simple, but the boundaries between them are where careers are protected or lost.

The green zone: permitted. These uses support your thinking without substituting for it. Using AI to search and map literature, in the Scout role from Chapter 8, provided every finding is verified before it enters your work. Polishing the language of paragraphs you wrote yourself: sharpening sentences, fixing grammar, improving flow. Formatting a reference list you have already verified. Brainstorming: asking what angles you might have missed, what a critic might say, how to structure an argument. In all of these, the ideas, the claims and the sources are yours; the machine assists their expression and their discovery.

The amber zone: permitted with disclosure. These uses shape the substance of the work, and honesty requires saying so. Substantial drafting assistance, where AI produces text that you then verify, rewrite and take responsibility for. Help with data analysis, such as generating analysis code or suggesting statistical approaches, which belongs in your methods description. Translation of your own writing between languages. The rule of the amber zone is simple: check your target journal's AI policy before submission, and disclose what you did in the manuscript, usually in the acknowledgements or methods. Journals differ in what they permit, but they are unanimous on one point: undisclosed use that the journal's policy requires you to disclose is misconduct, even if the use itself would have been acceptable.

The red zone: prohibited, always and everywhere. No journal policy, no deadline and no career pressure changes anything in this zone. Never cite a reference you have not verified and read. Never generate, alter or "improve" data; fabricated data is the deepest sin research has, whether a human or a machine fabricates it. Never list AI as an author, and never submit AI-generated text as your own analysis and argument; authorship means accountability, and a machine can carry none. Never use AI to write your peer reviews, and never upload someone else's unpublished manuscript into an AI tool; the manuscript sent to you for review was entrusted to you in confidence, and feeding it to a machine breaks that confidence regardless of how good the resulting review is.

One principle underlies all three zones, and it is the chapter's spine: **your name on the paper means you answer for every word, every number and every source in it.** "The AI did it" has never once succeeded as a defence, and it never will, because the accountability of authorship is precisely the thing that cannot be delegated. This is Chapter 1's two-question test at its maximum voltage: would you stand behind this with your name on it? In research, your name is already on it.

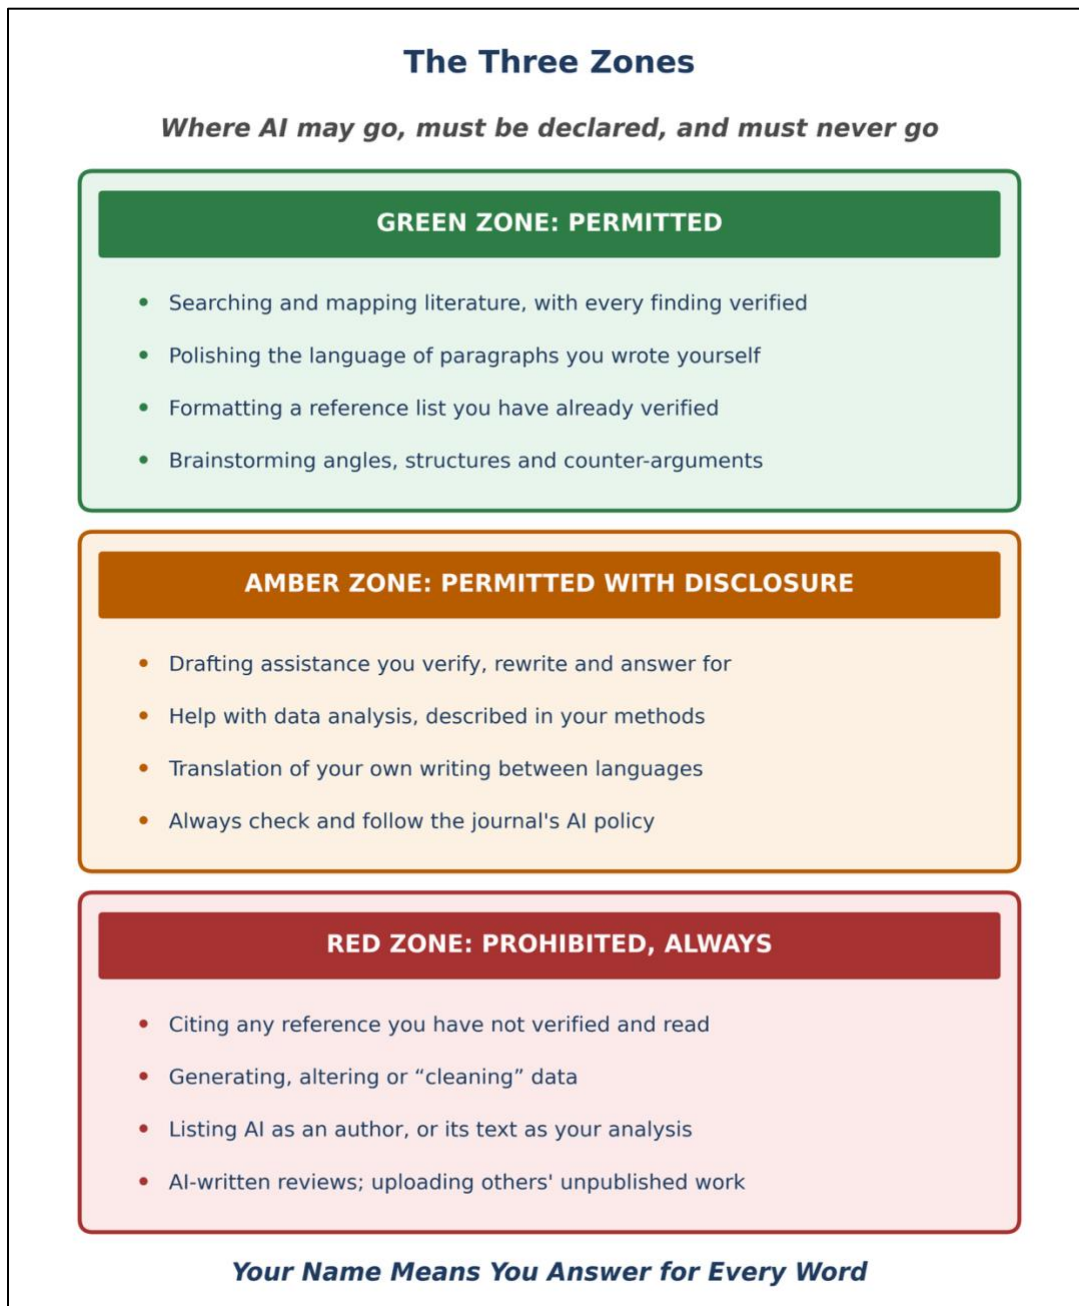


Figure 9: The Three Zones

Worked example: one paper, done honestly

The task: Dr Menon is writing a paper on student engagement in blended learning for a peer-reviewed journal.

Green zone work: she began with AI as Scout, asking for the main research directions on her topic, and received a useful map, including three suggested key papers. She searched

all three in Google Scholar: two existed and were genuinely relevant; the third did not exist at all. She was not surprised, only careful. Over two weeks, she found and actually read fourteen real papers, building her own reference list. When her draft was complete, she asked the AI to improve the clarity of her own paragraphs, accepting some suggestions and rejecting others.

Amber zone work: her statistics are honest, but her confidence in them is modest, so she asked AI to help her structure the analysis and generate the analysis code, which she then ran, checked and understood before trusting. Her methods section says so plainly, and her acknowledgements note that AI tools were used for language editing and analysis support, in line with the journal's stated policy, which she read before submitting.

The red line held: every one of her fourteen references has been opened and read. Her data went nowhere near an AI tool for "cleaning". The argument of the paper, its claims and its conclusions are hers, and in a viva on her own paper, she could defend every line, which is the test that matters.

The result: a paper that took slightly longer than a dishonest one and will outlast it. The reviewers' comments engaged with her actual argument. Nothing in the paper is waiting to be discovered.

Try this week

Run the verification test on your own field once, and you will never need convincing again. Ask an AI tool for ten key references on your specific research topic, with authors, years, journals and DOIs. Then check all ten in Google Scholar or Scopus. Typically, some will be real, some will be distorted versions of real papers, and at least one or two will be complete inventions, delivered in perfect confidence. The exercise costs twenty minutes and permanently changes how you treat every AI-suggested source: as a lead to investigate, never as a citation to use.

The takeaway

In research, AI can help you search, polish and think, but every reference must be read, every disclosure made, and every word answerable to your name. The green zone is generous, the amber zone is honest, and the red zone has no exceptions. The teacher still decides, and in research, the teacher also answers.

Chapter 10: Doing This Responsibly

The chapter that keeps the rest honest

Nine chapters have now said, in different ways, use it, plan with it, create with it, differentiate with it, let it help your feedback, your growth and your research. This final chapter says the other necessary thing: use it in a way you will still be proud of in ten years. Responsibility with AI is not one big dramatic decision. It is a handful of small habits, practised until they become reflexes, and this chapter names them.

The three quiet risks

The loud risks, cheating, fabricated references, and false accusations have had their chapters. Three quieter risks deserve naming before the book closes, because they arrive slowly and without alarms.

Privacy: A good working assumption for any AI tool is this: anything you type may be stored somewhere you do not control. That assumption settles most privacy questions instantly. Your students' names, marks and personal circumstances do not go in, as Chapter 6 established. Your colleagues' confidential matters do not go in. Other people's unpublished work does not go in, as Chapter 9 made absolute. Institutional documents that are not public do not go in. Describe situations without identifying people, and you lose almost nothing while protecting everyone.

Bias: AI tools learn from the world's existing text and inherit its habits: examples that default to Western settings, professions that quietly default to one gender, and scenarios that assume urban, affluent, English-speaking lives. None of this is deliberate, and all of it flows into your materials if nobody is watching. You are the somebody. When the AI's five examples of "a scientist" are all men, ask for better; when its case studies never leave America, bring them home, as the second filter in Chapter 3 taught. The machine reflects the world as it was written; the teacher represents the classroom as it actually is.

Over-reliance. This is the risk this book takes most seriously, because it is the one that arrives disguised as success. Every skill you delegate permanently is a skill that quietly weakens; the gym analogy from Chapter 7 applies to teachers exactly as it applies to students. The protection is simple and deliberate: keep some work unaided, by choice. Write some things from scratch. Do some marking arithmetic in your head. Draft the occasional difficult email yourself. Not because AI could not help, but because the muscles you keep are the judgment this whole book depends on. A teacher who can no longer do the task cannot supervise the machine doing it.

The framework: your personal AI policy

Institutions write AI policies, and yours may eventually hand you one. But the policy that will actually govern your daily choices is the one you write yourself, and it fits on a single page. Figure 10 gives you the template. It asks for four sentences, one for each risk this book has named.

My Personal AI Policy

Four sentences, written by you, kept where you work

I will never enter ...

students' names or work · confidences · others' unpublished manuscripts

.....

I will always verify ...

every fact · every answer key · every reference, in the database, myself

.....

I will still do unaided ...

the skills I refuse to let weaken; name them specifically

.....

I will always disclose ...

to journals as policy requires · to my students as I ask of them

.....

Signed

Date (revisit yearly)

A Policy in Your Own Words Governs Best

Figure 10: My Personal AI Policy

"I will never enter..." Your privacy line. Typically: students' names and work, colleagues' confidential information, others' unpublished manuscripts, non-public institutional documents.

"I will always verify..." Your accuracy line. Typically: every fact before it reaches students, every answer key before it marks anyone, every reference before it is cited.

"I will still do unaided..." Your skill line, and the one only you can write. Choose the abilities you refuse to let weaken, and name them specifically.

"I will always disclose..." Your honesty line. Typically: AI use in research per journal policy, and the same honesty about your own AI use that you ask of your students, because they will learn your real policy from what you do, not what you announce.

Write the four sentences, date the page, and keep it where you work. A policy you wrote yourself, in your own words, has a grip on behaviour that no circular from above will ever have. Revisit it once a year; it will need updating as tools change, and that is normal.

Worked example: one teacher's page

Dr Kapoor, who teaches management, wrote hers in ten minutes at the end of a faculty development programme, and it reads like this. *I will never enter: any student's name or work, anything a colleague told me in confidence, or any manuscript sent to me for review. I will always verify: every fact in my handouts, every answer key I use, and every reference I cite, in the database, myself. I will still do unaided: writing the first draft of anything that matters to me, and the final decision on every mark. I will always disclose: my AI use to journals as their policies require, and to my students when they ask me honestly, because I ask the same of them.* She signed it, dated it, and taped it inside her cupboard door. She reports that the third line is the one she rereads most.

Where do you stand now?

In Chapter 1, you scored yourself on ten statements. Here they are again, unchanged. Answer honestly, without looking back at your earlier answers first: **Not yet = 1, Occasionally = 2, Regularly = 3, Confidently = 4.**

#	Statement	Not yet (1)	Occasionally (2)	Regularly (3)	Confidently (4)	My score now	My score in Chapter 1
1	I can describe in one sentence	1	2	3	4		

	what tools like ChatGPT, Gemini or Claude actually do.						
2	I can write a clear instruction to an AI tool and improve it when the first result disappoints.	1	2	3	4		
3	I check AI-generated content for errors before using any part of it.	1	2	3	4		
4	I can spot when an AI answer sounds confident but is actually wrong or invented.	1	2	3	4		
5	I have used AI to help prepare teaching materials such as questions, examples or outlines.	1	2	3	4		
6	I have thought about how AI changes the way I should design assessments.	1	2	3	4		

7	I have spoken with my students about acceptable and unacceptable uses of AI.	1	2	3	4		
8	I know which uses of AI in research are legitimate and which cross ethical lines.	1	2	3	4		
9	I am careful about what personal or student information I enter into AI tools.	1	2	3	4		
10	I can explain and defend any AI-assisted work I produce, with my name on it.	1	2	3	4		
	Total score					/40	/40

Now compare the two totals. If your score rose, the distance between the numbers is learning you can name, item by item; notice which statements moved most, because those are the chapters that worked on you. If a score fell anywhere, that is progress of a different kind: you now know enough to judge yourself more honestly than you could in Chapter 1, and honest self-assessment is itself one of this book's quiet goals. Either way, the ten statements are yours to keep. Return to them in six months, because the tools will have changed by then, and so will you.

Try this week

Two actions close the book. First, write your personal AI policy, four sentences, today, while the chapters are fresh, and put it where you will see it. Second, choose the one chapter of this book that most challenged your current practice, and do its "try this week" again,

properly this time. A book like this succeeds not when it is finished but when it is used, and only you can decide that part.

The takeaway

Protect what is private, correct what is biased, and keep the skills that make you the judge. Ten chapters have shown what AI can do for a teacher; not one of them found something it should decide. It drafts; you decide; it suggests; you select; it assists; you answer. That was the whole book, and it fits in four words. The teacher still decides.

A Closing Note: The Decision That Cannot Be Delegated

You have reached the end of a short book with one long argument, and you already know what it is. Artificial intelligence can draft your plans, generate your materials, phrase your feedback and search your literature, and this book has encouraged you to let it. But at every point where quality is judged, where fairness is weighed, where a mark is given, or a claim is made under your name, the decision has remained where it always belonged. The tools will keep changing, and faster than any book can follow. The principle will not.

If the chapters have done their work, you now have more than an opinion about AI; you have a set of habits waiting to be practised. Ten small actions, one per chapter, are gathered in Appendix A. Do not attempt them all at once. Once a week is enough, and by the end of a term, the teacher who began this book will be quietly different from the one finishing it. Your two self-assessment scores already know this.

This book is published with open access, deliberately. Knowledge about teaching well should not wait behind a paywall, least of all now. Share it freely with colleagues, discuss it in your faculty rooms, use it in your development programmes, and adapt what works to your own context; the licence asks only that you credit the source and keep it non-commercial. A companion guide for teacher educators and pre-service teachers already exists, and guides for students and for academic leaders will follow, because every member of this ecosystem now needs the same thing this book offered you: a way to use AI every day without losing your judgement.

Teaching has survived the printing press, the calculator, the internet and the smartphone, and it will survive this too, changed but recognisable, for one reason above all the others. Machines have never once known what a student needs. Teachers do. The teacher still decides.

About the Author



Adit Gupta, Professor and Principal, MIER College of Education (Autonomous), Jammu

Adit Gupta is Professor and Principal of MIER College of Education (Autonomous), Jammu, and Director of the MIER Group of Institutions. He holds a PhD in Learning Environments from Curtin University, Australia and brings nearly three decades of experience in teacher education, educational technology, and institutional leadership.

His research spans AI in education, psychometric instrument development, blended learning, and technology-enhanced teacher preparation. He has developed a portfolio of original frameworks and instruments for understanding how educators and institutions engage with emerging technologies, including the AI Interaction Knowledge framework, which runs throughout this guide.

Between 2021 and 2024, he led a three-phase technology-enhanced learning transformation project at MIER College of Education in partnership with the Commonwealth of Learning, resulting in the development of over twenty blended courses and a benchmarking score of 4.1 out of 5 across ten domains of institutional TEL practice. That project is the direct institutional foundation from which this guide grew.

He serves as Regional Editor of Learning Environments Research, published by Springer, and as Editor-in-Chief of the MIER Journal of Educational Studies, Trends and Practices. His work has been recognised through the Australian Government Endeavour Executive Award and the United States Government International Visitor Leadership Programme.

This guide follows his companion volume, *AI in the Classroom Before the Classroom* (2026), and continues the same mission: something a teacher in Jammu, Jaipur, or Jorhat can pick up on a Sunday evening and actually use on Monday morning. He hopes it is useful. He would be glad to hear from anyone for whom it was.

adit@mier.in

MIER College of Education (Autonomous), Jammu

Appendix A: My AI Practice Plan

This book asks you to try one small thing at the end of every chapter. Here they all are, gathered on one page. Do not attempt them all at once; once a week is enough, and a term is plenty. Choose your starting point, write a target date against each action, and tick them off as you go. By the last tick, your practice will have changed more than any workshop could have changed it.

Chapter	The action	My Target Date	Done
1. Before You Begin	Give an AI tool one routine task, examine its output, find its errors, and apply the two-question test.		
2. Planning and Preparation	Write a full five-sentence brief for one class, ask for two plan options, and change three things while shaping.		
3. Creating Teaching Materials	Generate five multiple-choice questions with answers, attempt all five yourself, and check the answer keys.		
4. Teaching the Large Class	Take your hardest concept and prepare it in three versions: simple, standard and stretch. Teach the simple version first.		
5. Assessment in the Age of AI	Ask AI to complete your most vulnerable assessment. If it would pass, redesign the task with one move.		
6. Feedback at Scale	Tally the repeated errors in one set of marking, and turn the tally into a whole-class feedback sheet.		
7. Your Students Are Already Using It	Give your next assignment a traffic-light colour, with one sentence of reasoning, and explain it in class.		
8. Your Own Growth	Spend twenty minutes asking AI about the one concept you have pretended to understand for years.		
9. AI in Your Research and Publication	Ask AI for ten references on your research topic, then verify all ten in a real database. Count the inventions.		
10. Doing This Responsibly	Write your four-sentence personal AI policy, sign it, date it, and keep it where you work.		

When the last box is ticked, return to the self-assessment at the end of Chapter 10 and score yourself a third time. Then pass this book to a colleague; it is open access, and it was written to travel.

Appendix B: The Prompt Bank

Every worked example in this book used a prompt, and they are collected here as reusable templates, one per chapter. Replace the square brackets with your own subject, students and situation; the brackets are where your context goes, and Chapter 2 explained why the context is what makes the output worth having. Each prompt is a starting point, not a script. Say more than the template where you can, and remember that everything that comes back passes through your filters before it goes anywhere near a student.

Chapter 1: The routine task (Dr Sharma's notice)

"Write a short, polite notice for [your students] informing them that [the change and its reason]. Mention that [what stays the same]. Formal but warm tone, under 100 words."

Chapter 2: The briefed lesson plan (Dr Verma's memory lecture)

"I teach [subject] to [year and programme] students in India. I need a plan for one [length] class on [topic] for about [number] students with [their characteristics]. [State your room and resource constraints.] The prescribed textbook covers [scope]. I want at least one activity that gets students participating without needing any materials. Give me two different plan options with rough timings."

Chapter 3: The local case study with questions (Mr Qureshi's vegetable seller)

"Write a one-page case study for [year and programme] students in India about [concept], based on [a local, familiar situation]. Simple English, short sentences, prices in rupees. Then create five multiple-choice questions on the case, with four options each, and indicate the correct answers."

Remember: attempt every question yourself before trusting the answer key.

Chapter 4: The three versions (Dr Bhat's standard deviation)

"I teach [subject] to [students], many of whom [their main difficulty]. Explain [concept] in three versions: first, a simple version using an everyday analogy and no [formula/jargon]; second, a standard textbook version with one small worked example; third, a stretch version posing a thought question that applies the idea to [their field]. Simple English throughout."

Chapter 5: The redesigned assessment (Dr Nair's term paper)

"Here is an assignment I have used for years: [paste it]. Redesign it so that students must collect their own local material [interviews, observations or data], submit staged work at three points in the term, and present briefly in class. Keep the same topic and learning goals."

Give me the new task description, the three checkpoint requirements, and a simple marking rubric."

Chapter 6: The whole-class feedback sheet (Dr Joshi's four errors)

"I teach [subject] in India. My class of [number] made these errors in their [assignment]: [your observations, in your own words]. Draft a one-page whole-class feedback sheet. For each error: name it plainly, explain why it weakens the work, and show a two-line corrected example using a generic case. Warm, encouraging tone. Simple English."

Remember: your observations go in; no student's work or name ever does.

Chapter 7: The traffic-light course outline (Dr Fernandes's first day)

"Here are the assessment tasks for my course: [list them]. For each task, I have decided its AI rule: [red, amber or green, with your reasoning]. Draft the AI-use section of my course outline: state each task's colour and reason in one sentence and include a two-line sample disclosure note that students can copy for amber tasks. Plain, friendly language."

Chapter 8: The patient tutor (Mr Thomas's chi-square)

"Explain [the concept you have avoided] to someone who has stayed away from [the field] for twenty years. No [formulas/jargon] yet. I will ask follow-up questions."
And a week later, the Coach version: *"I will now explain [the concept] back to you in my own words: [your explanation]. Point out every error and gap. Do not be polite."*

Chapter 9: The verification test (the exercise that ends all complacency)

"Give me ten key references on [your specific research topic], with authors, years, journals and DOIs."

Then verify all ten in Google Scholar or Scopus and count the inventions. This prompt exists to be distrusted; that is its lesson.

Chapter 10: The policy stress test (after writing your four sentences)

"Here is my personal AI policy: [paste your four sentences]. Describe five realistic situations in my working week that would test this policy, including at least one where the tempting choice and the right choice differ. Do not soften them."

Appendix C: Quick Decision Card

The Quick Decision Card

One page, three tools · pin it where you work

BEFORE ANY AI TASK: THE TWO-QUESTION TEST

1. Am I still doing the thinking?

2. Would I stand behind this with my name on it?

Both yes: proceed, the work is yours · Either no: stop and reclaim the task

SETTING STUDENT TASKS: THE TRAFFIC LIGHT



RED: AI not permitted – this task builds a skill you must own



AMBER: permitted with a short disclosure note, never punished



GREEN: AI use expected – your judgement of it is what is marked

IN RESEARCH: THE THREE ZONES



GREEN: search, polish and format – with every finding verified



AMBER: drafting and analysis help – disclosed per journal policy



RED: unverified citations, altered data, AI reviews – never

It drafts, you decide · it suggests, you select · it assists, you answer

THE TEACHER STILL DECIDES

USING AI EVERY DAY WITHOUT LOSING YOUR JUDGEMENT.

Everyone has an opinion about AI in higher education. This short guide offers a working method instead: ten brief chapters on planning, teaching, assessing and publishing with AI as an assistant that never becomes the judge. Every chapter ends with one action that costs nothing and can be tried this week.

Adit Gupta is Professor and Principal of MIER College of Education (Autonomous), Jammu.



OPEN ACCESS · CC BY-NC 4.0 · OAR.MIERCOLLEGE.IN
ISBN 978-81-952808-4-1