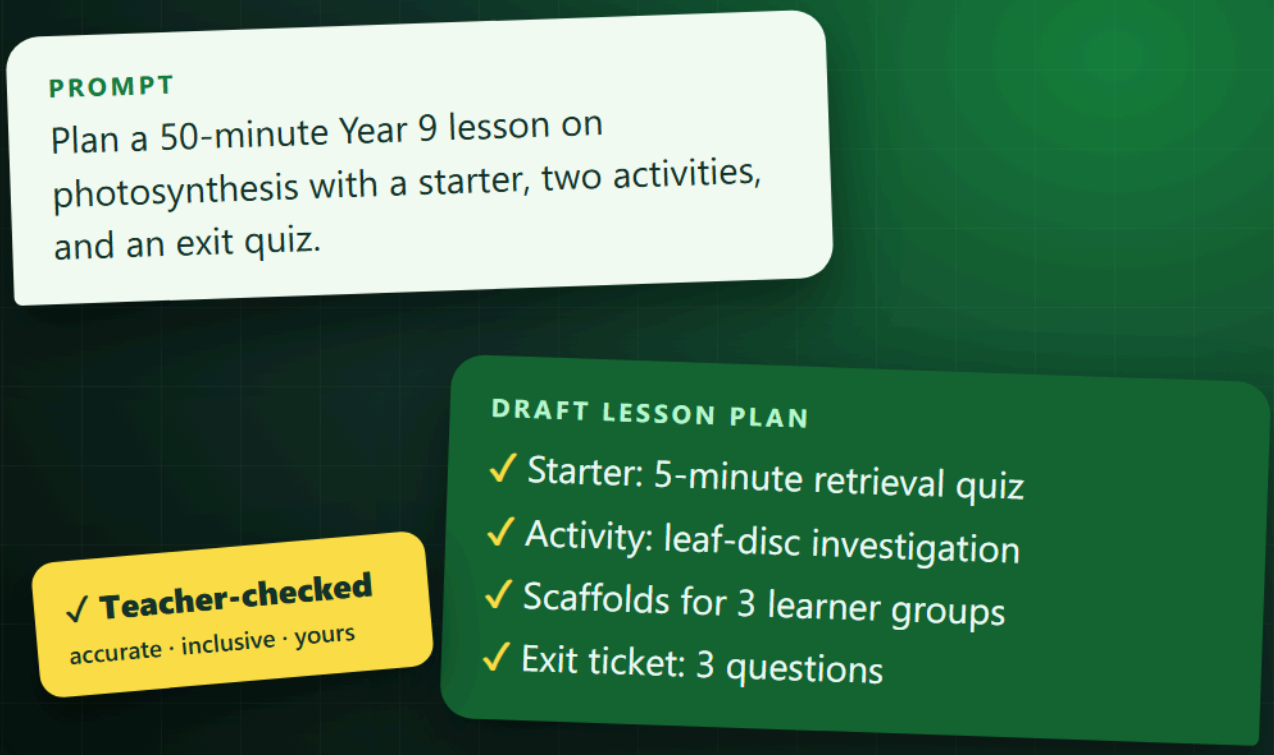


THE PRACTICAL GUIDE

AI for Teachers and Trainers

Using Generative AI for Planning, Teaching, Assessment, and Feedback

Prompting · Lesson plans · Differentiation · Quizzes & rubrics
Feedback · Integrity · AI literacy · Workplace training

PLAN · CREATE · ASSESS · SUPPORT · LEAD

THENEXBOOKS

AI for Teachers and Trainers

A Practical Guide to Using Generative AI for Planning, Teaching, Assessment, and Feedback

First Edition

THENEXBOOKS

AI for Teachers and Trainers
A Practical Guide to Using Generative AI for Planning, Teaching, Assessment, and Feedback

First Edition, 2026

Copyright © 2026 thenexbooks (thenexbooks.gumroad.com). All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in reviews and certain other non-commercial uses permitted by copyright law. Prompts, templates, and checklists may be used and adapted freely in your own teaching and training.

The information in this book is provided "as is" without warranty of any kind. AI tools, their features, and their terms of use change frequently; always check the current terms of any tool, your institution's policies, and the laws that apply where you work. Nothing in this book is legal advice. The publisher assumes no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

Brightwater Secondary School, Brightwater Health, Brightwater Health Academy, and all people and events associated with them are fictional. Product names such as ChatGPT, Claude, Gemini, and Microsoft Copilot are trademarks of their respective owners and are used for identification only.

Preface

Generative AI arrived in classrooms and training rooms faster than any technology before it. Learners started using it before most institutions had a policy. Colleagues began sharing prompts in staff rooms and team chats. Some educators felt excited, some anxious, and many simply exhausted by one more thing to learn.

This book is written for all of them. It does not assume you are a technology enthusiast, and it does not ask you to hand your teaching over to a machine. It starts from what you already know — how people learn, what good lessons and training sessions look like, what makes feedback useful, and why relationships matter — and shows where AI can genuinely help: planning faster, creating better-differentiated materials, giving more feedback to more learners, and reclaiming time for the parts of the job only a human can do.

It is equally honest about the limits. AI tools make confident mistakes, reflect biases in their training data, raise real privacy questions, and can tempt learners to skip the thinking that learning requires. Every chapter shows you how to check, adapt, and decide — because the professional judgment stays with you.

Who This Book Is For

- **School teachers** in primary and secondary education who want practical, classroom-ready ideas.
- **College and university instructors** rethinking assignments, feedback, and integrity.
- **Corporate, vocational, and healthcare trainers** who design and deliver workshops, onboarding, and compliance training.
- **Instructional designers** building courses, e-learning, and blended programs.
- **Education and L&D leaders** writing policy, supporting staff, and measuring impact.

No programming or technical background is needed. Every technical term is explained the first time it appears, and the glossary in Appendix D collects them all.

What You Will Learn

- How generative AI works in plain language — and why it can be confidently wrong.
- How to choose tools, protect learner privacy, and write prompts that produce useful results.
- How to plan courses, lessons, and training sessions, and create accurate, inclusive materials.
- How to redesign assessment, build quizzes and rubrics, give feedback at scale, and handle academic integrity fairly.
- How to teach AI literacy and use AI for tutoring, role-play, inquiry, and workplace coaching.
- How to save time on administration, build reusable assistants, measure impact, lead adoption, and protect your wellbeing.

- Two complete case studies you can adapt to your own school, college, or organization.

How This Book Is Organized

Part I — Foundations (Chapters 1–5) explains what AI is, how to choose tools, how to prompt, and how to use AI responsibly. **Part II — Planning and Creating Learning** (Chapters 6–10) covers course and lesson design, learning materials, differentiation, and multimedia. **Part III — Assessment and Feedback** (Chapters 11–15) covers assessment design, quizzes and rubrics, feedback, integrity, and grading. **Part IV — AI in the Learning Experience** (Chapters 16–20) puts AI in learners' hands wisely. **Part V — Productivity, Leadership, and Growth** (Chapters 21–26) covers administration, custom assistants, online learning, impact, leadership, and wellbeing. **Part VI — Case Studies** (Chapters 27–28) brings everything together. The appendices provide a prompt cheat sheet, templates, checklists, and a glossary.

You do not need to read the book in order. Chapters 1, 4, and 5 are the essential foundation; after that, go straight to the chapters that match your role and your most pressing needs.

A Note on Tools

This book is **tool-neutral**. The prompts and workflows work with any capable AI chat assistant, whether a general tool or one built into your school's or company's systems. We deliberately avoid step-by-step instructions for specific products, because features, prices, age requirements, and privacy terms change often. Whenever a decision depends on a tool's current terms, we tell you what to check.

The Running Example: Brightwater

Throughout the book we follow two fictional settings. At **Brightwater Secondary School**, science teacher Lena Ortiz plans units and supports a diverse Year 9 class, English and history teacher Sam Adeyemi rethinks essay writing and feedback, learning-support lead Mei Tanaka makes materials accessible, and deputy principal Grace Holloway leads the school's AI policy. At **Brightwater Health Academy**, the training department of a regional healthcare provider, trainer Jordan Reyes runs onboarding and patient-communication workshops, instructional designer Priya Nair builds e-learning, and L&D manager Tom Becker needs to show that training makes a difference. Their challenges are the ones you face every week.

Conventions Used in This Book

> **TIP**
>
> A recommended practice or a technique that saves time or improves results.

> **NOTE**
> Background information or a useful aside.

> **WARNING**
> A common mistake that could affect learners, privacy, or accuracy.

Prompts and sample AI outputs appear in shaded boxes. Placeholders in square brackets, such as `[TOPIC]` or `[YEAR GROUP]`, are for you to replace. Sample outputs are illustrative and shortened; your tool's responses will differ. Numbers labeled *hypothetical* or *illustrative* are examples, not research findings.

Contents

Preface

PART I — FOUNDATIONS

1 Why AI Matters for Teaching and Training

2 How Generative AI Works — in Plain Language

3 Choosing and Setting Up AI Tools

4 Prompting Fundamentals for Educators

5 Responsible Use: Privacy, Safety, Bias, and Policy

PART II — PLANNING AND CREATING LEARNING

6 Curriculum and Course Design with AI

7 Lesson Plans and Training Sessions

8 Creating Learning Materials: Readings, Slides, Handouts, and Examples

9 Differentiation, Accessibility, and Inclusive Learning

10 Visuals, Multimedia, and Language Learning

PART III — ASSESSMENT AND FEEDBACK

11 Designing Assessment in the Age of AI

12 Question Banks, Quizzes, and Rubrics

13 Feedback at Scale — Without Losing the Human Touch

14 Academic Integrity and AI-Resilient Assignments

15 Grading, Moderation, and Learning Data

PART IV — AI IN THE LEARNING EXPERIENCE

16 Teaching AI Literacy to Learners

17 AI Tutors and Study Assistants

18 Discussion, Debate, Simulations, and Role-Play

19 Projects, Inquiry, and Critical Thinking with AI

20 AI in Workplace Training and Coaching

PART V — PRODUCTIVITY, LEADERSHIP, AND GROWTH

21 Saving Time on Administration and Communication

22 Building Custom Assistants and Prompt Libraries

23 AI for Online, Hybrid, and Blended Learning

24 Measuring Impact: Evaluating AI in Your Classroom or Program

25 Leading AI Adoption: Policy, Professional Development, and Change

26 Wellbeing, Ethics, and the Future of Teaching

PART VI — CASE STUDIES

27 Case Study: Brightwater Secondary's Science Department

28 Case Study: Brightwater Health Academy's Onboarding Program

APPENDICES

A Prompt Cheat Sheet

B Templates

C Checklists

D Glossary

PART I

Foundations

What generative AI is, how it works in plain language, how to choose tools, how to write prompts that get useful results, and how to use AI responsibly with learners.

1 · Why AI Matters for Teaching and Training
2 · How Generative AI Works
3 · Choosing and Setting Up AI Tools
4 · Prompting Fundamentals for Educators
5 · Responsible Use

CHAPTER 1 · PART I — FOUNDATIONS

Why AI Matters for Teaching and Training

Generative AI will not teach your learners for you, but it can take a surprising amount of drafting, adapting, and organizing off your desk. This chapter sets realistic expectations, addresses the fears honestly, walks you through a safe first ten-minute try, and introduces the two fictional organizations whose people will accompany you through the rest of the book.

A New Kind of Assistant in the Staffroom

Every educator knows the feeling. It is late on a Sunday evening. You have a lesson to plan for Monday, a stack of work to respond to, a parent or manager email you have been putting off, and a nagging sense that the worksheet you made three years ago no longer fits the group in front of you. The teaching itself — the moment when a learner finally sees the idea — is the part you love. Much of the rest is preparation, adaptation, and administration that expands to fill whatever time you give it.

Generative AI tools arrived in this landscape quickly. Within a short span, chat assistants that can draft text, explain ideas, summarize documents, and answer follow-up questions became available to almost anyone with a web browser. Widely known examples include ChatGPT, Claude, Gemini, and Microsoft Copilot, and many education platforms and office tools now include similar features. These products change constantly — their features, prices, and terms shift from month to month — so in this book we refer to "your AI assistant" or "a chat assistant" and focus on skills that transfer from one tool to another.

The word **generative** matters. Earlier kinds of software in education mostly *stored* or *sorted* things: gradebooks, learning management systems, quiz engines. Generative AI *produces* new text (and, in some tools, images or audio) in response to instructions written in ordinary language. You type a request such as "Draft five discussion questions about the causes of the First World War for 15-year-olds," and within seconds you have a draft to work with. That is genuinely new, and it is why so many educators are curious, excited, and uneasy all at once.

This book is written for all of those feelings. It is for school teachers, college and university instructors, corporate and vocational trainers, instructional designers, and education leaders. You do not need any programming knowledge. You need only what you already have: knowledge of your subject, knowledge of your learners, and a professional commitment to doing right by them. AI tools are most useful in the hands of people who have exactly those things.

Meet Brightwater: Our Running Example

Advice about AI can feel abstract, so throughout this book we follow the people of two fictional organizations that share a town and a name. One is a school; the other is a workplace training department. Between them, they cover most of the situations our readers face. None of these people or organizations is real, but their challenges will probably sound familiar.

Brightwater Secondary School

Brightwater Secondary is a public secondary school serving students aged 11 to 18. It is a busy, ordinary school: large classes, a wide range of needs, a timetable that never has quite enough slack in it, and staff who care a great deal. Four members of staff will appear again and again.

Lena Ortiz teaches science. Her Year 9 class, 9C, has 28 students aged 13 to 14. Three of them are English-language learners at different stages of fluency, and two have dyslexia. Lena is a strong, organized teacher who likes practical work, but she spends many evenings adapting materials so that every student in 9C can access the same core ideas. Her current topic is photosynthesis, and we will watch her plan, differentiate, and assess it with AI support in several chapters. Lena's question is simple: "Can this give me back some of my evenings without making my lessons worse?"

Sam Adeyemi teaches English and history and works with Year 11 students (aged 15 to 16) on essay writing. Sam is the most skeptical member of staff we will meet. Sam's worry is that students will use AI to write their essays and that the craft of building an argument — the thing Sam most wants to teach — will quietly disappear. Sam's journey through this book is not from skeptic to enthusiast; it is from skeptic to *informed* skeptic, someone who knows exactly where AI helps, where it harms, and how to design assignments accordingly.

Mei Tanaka is head of learning support, responsible for students with special educational needs and for accessibility across the school. Mei sees AI through two lenses at once. On one hand, tools that can rewrite a text at a simpler reading level, generate a glossary, or restructure instructions into short steps could be enormously helpful. On the other, Mei knows that students with additional needs are often the most affected when technology is introduced carelessly, and that their personal information is especially sensitive.

Grace Holloway is the deputy principal leading the school's AI policy. Grace's job is to make sure that whatever the school does with AI is safe, lawful, fair, and educationally worthwhile — and that staff are supported rather than simply handed a rulebook. Grace has to balance enthusiastic early adopters, anxious colleagues, cautious parents, and the practical realities of the school's approved software list.

Brightwater Health Academy

Brightwater Health Academy is the training department of Brightwater Health, a fictional regional healthcare provider. The Academy trains new hires, runs refresher courses, and builds e-learning modules for staff across clinics and offices. Its learners are adults with jobs to do, limited time, and a

strong interest in training that is directly useful.

Jordan Reyes is a trainer who runs new-hire onboarding and customer-service workshops focused on patient communication — how front-desk staff, call-center colleagues, and new clinical support workers speak with patients and families, especially when people are worried, confused, or upset. Jordan is energetic and practical and wants realistic scenarios and role-plays that do not feel stale after the fifth cohort.

Priya Nair is an instructional designer who builds e-learning modules on topics such as data-privacy compliance and infection control basics. Priya works with subject-matter experts, writes scripts and knowledge checks, and cares about accuracy above almost everything else, because a mistake in a compliance or infection control module can have real consequences.

Tom Becker is the learning and development (L&D) manager. Tom must show leadership that training makes a difference, manage a budget, and make procurement decisions about new tools. Tom asks the questions every manager should ask about AI: What does it cost? Where does our data go? How will we know it is helping?

> **NOTE**
>
> Brightwater Secondary School, Brightwater Health Academy, and everyone who works there are fictional. We use them so that examples feel concrete and continuous. When you see Lena planning a lesson or Jordan designing a role-play, imagine your own version of that task.

What Generative AI Can Do for Educators

Let us be specific. When used well, a chat assistant is good at a particular family of tasks. Most of them involve producing a *first draft* or a *variation* of something that you, the professional, then review and shape.

Task type	School example	Training example
First drafts	Lesson outline	Workshop agenda
Variations	Three reading levels	Scenarios by role
Question writing	Quiz items	Knowledge checks
Explanations	Analogies for a concept	Plain-language policy
Reorganizing	Notes into a handout	SME notes into a script
Communication	Newsletter draft	Course reminder email
Brainstorming	Hook activities	Role-play ideas
Feedback support	Comment bank ideas	Observation checklist

Notice the pattern. In every row, the AI produces raw material. It does not decide whether the material is right for your learners, whether it is accurate, or whether it fits your curriculum or organization. Those decisions stay with you. The time saving comes from not starting with a blank page, and from how quickly you can ask for alternatives: "Make it shorter." "Now a version for learners who are new to the topic." "Give me three more examples, each from a different workplace setting."

For Lena, a realistic early win might be asking for a set of retrieval-practice questions on last week's topic, then choosing and editing the best eight. For Jordan, it might be generating ten short patient-communication scenarios with different emotional tones, then selecting four that match what front-desk staff actually encounter. Neither of them hands over their judgment. Both of them skip the slowest part of the job: getting something onto the page.

What Generative AI Cannot Do (and Why It Matters)

An honest picture of limitations is just as important as the list of strengths. Chapter 2 explains *why* these limitations exist; here we simply name them.

- **It can be confidently wrong.** AI tools sometimes produce statements that sound authoritative but are false, including invented references, incorrect dates, or subtly wrong explanations. This is often called **hallucination**. Everything factual must be checked, especially in science, health, law, and history.
- **It does not know your learners.** The tool has never met 9C or Jordan's latest cohort. It cannot see who is anxious, who is bored, who has just had a difficult week, or who needs a word of encouragement. It only knows what you tell it — and you should be careful what you tell it.
- **It does not know your context by default.** It may not know your curriculum, your exam board, your organization's policies, or the local rules that apply to you, unless you provide that information. Even then, it can misapply it.
- **It does not take responsibility.** If a worksheet contains an error, or feedback is unfair, the professional who used it is accountable. AI can draft; it cannot be answerable.
- **It is not neutral.** Because these tools learn from large collections of human-written text, they can reflect the biases in that text — about gender, culture, disability, language, and more. You need to look for this.
- **It does not build relationships.** Trust, belonging, motivation, and the sense that someone believes in you are built between people. That is the heart of teaching and training, and it remains human work.

> **WARNING**
>
> Never paste real learners' names, grades, health information, or other personal details into an AI tool unless your institution has explicitly approved that tool for that purpose. Throughout this book, examples use anonymized or invented details. Chapter 5 covers privacy in depth; Chapter 3 shows you how to check a tool's data settings.

Where AI Saves Time — and Where Judgment Is Essential

A useful way to think about any task is to ask two questions. First: *how much of this task is producing text or structure?* Second: *how much depends on professional judgment about specific people?* Tasks high on the first and low on the second are where AI tends to help most. Tasks high on the second need you in the driver's seat, even if AI supports some steps.

Task	AI help	Your judgment
Drafting quiz questions	High	Check accuracy, level
Adapting reading level	High	Keep meaning intact
Writing a course email	High	Tone, facts, dates
Designing a unit sequence	Medium	Curriculum fit
Feedback on learner work	Medium to low	Fairness, relationship
Final grades and decisions	Low	Fully human
Pastoral or welfare issues	Very low	Fully human

The table is a starting point, not a rule. Feedback is a good example of a gray area. AI can help you write clearer comments, build a bank of common feedback phrases, or suggest next steps for a typical misconception. But deciding what a particular learner needs to hear, and how to say it so that they keep going, is a human judgment. Chapter 13 explores this balance in detail.

High-stakes decisions — final grades, disciplinary matters, certification, safeguarding concerns, performance management in the workplace — should never be delegated to an AI tool. Most institutional policies already say this, and so does good sense. Where AI is used anywhere near such decisions, a qualified human should make the decision and be able to explain it.

The Educator as Designer, Not Replaced

A helpful mental shift is to think of yourself as the **designer** of learning and the editor of AI output, rather than as someone competing with a machine. Architects use software to draft plans, but the building still reflects the architect's understanding of the site, the client, and the safety codes. In the same way, you decide what learners need to learn, what evidence will show that they have learned it, and what experiences will get them there. AI can help you produce the materials faster and in more variations.

This framing connects to ideas many educators already use. **Backward design**, associated with Grant Wiggins and Jay McTighe, asks you to start with the desired results, then decide on acceptable evidence, and only then plan learning activities. AI is useful at the third stage and sometimes the second, but the first — deciding what matters — is yours. Bloom's taxonomy (in its 2001 revised form) gives you language to ask for tasks at different cognitive levels, from remembering to creating.

Universal Design for Learning, developed by CAST, encourages multiple means of engagement, representation, and action and expression; AI can help you produce those multiple means more quickly. We return to all of these frameworks in later chapters.

Consider how the designer role plays out for Sam. Sam does not want AI writing essays, but Sam could use AI to generate a deliberately flawed model paragraph for students to critique, or to produce three different thesis statements on the same question so the class can rank them and explain why. In both cases, Sam is designing a learning experience, and the AI output is a teaching object — something to analyze — rather than a replacement for student thinking.

Or consider Priya. An infection control module needs accurate content signed off by a qualified subject-matter expert. AI does not change that. What it can change is how quickly Priya turns expert notes into a draft script, a set of scenario-based questions, and a plain-language summary — all of which then go back to the expert for review. Priya becomes the designer and editor of a faster process, not a passive recipient of machine output.

> **TIP**
>
> A simple rule of thumb: **AI drafts, you decide.** If you cannot confidently check an output yourself (because it is outside your expertise), either find someone who can or do not use it.

Realistic Expectations

Many educators have one of two early experiences with AI. Some are dazzled: the first output looks polished and complete, and it seems as though planning has been solved. Others are disappointed: the output is generic, bland, or subtly wrong, and they conclude that the whole thing is hype. Both reactions are understandable, and both miss the middle ground where most of the value lives.

Here are expectations that tend to hold up over time.

- **The first answer is rarely the best answer.** Useful results usually come from a short conversation: you ask, review, and refine. Chapter 4 teaches this process step by step.
- **Specific requests get specific results.** A vague request ("make a worksheet on fractions") gets a generic worksheet. A request that describes your learners, their prior knowledge, and your goal gets something much closer to usable.
- **Editing is part of the job.** Expect to spend real time reviewing and adjusting. The saving is in the drafting, not in skipping quality control.
- **Time savings build gradually.** The first few attempts may take longer than doing the task yourself, because you are learning a new skill. Once you have a few reliable prompts, the savings become more noticeable.
- **Some tasks are not worth it.** If you can write a two-line email in thirty seconds, you do not need AI. Save it for tasks where the drafting burden is genuinely heavy.

We do not quote percentages of time saved in this book, because the honest answer is that it depends — on the task, the tool, your experience, and how carefully you check the results. What we can say is that many educators find specific, repeatable tasks where AI makes a clear difference, and that finding those tasks is a skill you can develop.

Common Fears, Addressed Honestly

It would be strange if you felt no unease about AI. Educators have seen many technologies arrive with bold promises. Some fears about AI are exaggerated; some are well founded. Here we take the three most common ones seriously.

"Students will just cheat."

This is Sam's first fear, and it deserves a straight answer: yes, some learners will use AI to produce work they present as their own, just as some learners copied from friends, the internet, or essay-writing services before AI existed. AI makes this easier and faster, and it affects some kinds of assessment — especially unsupervised take-home essays and short written answers — much more than others.

It is also true that the most common technical "solution," AI-text detection software, is unreliable. Detectors can produce false positives, flagging genuine human writing as AI-generated, and there is particular concern that English-language learners and others whose writing follows simpler or more formulaic patterns may be wrongly flagged. A detector score should never be the sole basis for an academic integrity decision.

The more durable response is design: assessments that value process as well as product, in-class writing, oral explanation, drafts and reflections, personal or local connections, and clear guidance about when AI use is and is not permitted. Chapters 11 and 14 are devoted to this. For now, it is enough to know that the problem is real, that it is manageable, and that the answer is mostly pedagogical rather than technical.

"AI will take my job."

No one can predict the long-term future of any profession with certainty, and we will not pretend otherwise. What we can say is that the core work of teaching and training — understanding learners, building relationships, motivating, judging what someone needs next, maintaining a safe and purposeful environment, and taking professional responsibility — is precisely the work that AI tools do poorly. Learners, parents, employers, and regulators continue to expect a qualified, accountable human to be responsible for learning.

What is more likely to change is the *shape* of the job. Some drafting and administrative tasks may shrink. Skills such as evaluating AI output, designing AI-resilient assessments, and teaching learners to use AI responsibly may grow. Educators who understand these tools well are in a stronger position

to influence how they are used in their institutions — which is one reason Grace and Tom both want their colleagues to read books like this one.

"We'll lose our skills — and so will learners."

This concern, sometimes called **deskilling**, is thoughtful and important. If a new teacher never writes a lesson plan from scratch, will they develop the planning judgment that experienced teachers have? If learners outsource every first draft, will they learn to write? If trainers rely on generated scenarios, will they lose touch with what their learners actually face?

These are real risks, and the answer is intentional use. For your own practice, keep doing some tasks without AI, especially those at the heart of your expertise, and use AI where it frees time for higher-value work. For learners, distinguish between tasks where the *struggle is the point* (learning to structure an argument, working through a calculation) and tasks where AI support is reasonable (checking a draft's clarity, generating practice questions). Chapter 16 covers teaching AI literacy, including when *not* to use AI.

> **NOTE**
>
> International bodies have also engaged with these questions. UNESCO's 2023 publication "Guidance for generative AI in education and research" discusses a human-centered approach, age-appropriate use, and the need for institutions to validate tools before adopting them. It is worth reading alongside your own local policies.

Your First 10-Minute Try

Reading about AI is useful, but trying it is better. This short exercise is designed to be safe, low-stakes, and genuinely informative. It uses no personal data, involves no learners, and takes about ten minutes.

Before you start

1. Check your institution's policy. If your school or organization has an approved AI tool, use that one. If you are unsure, ask — Grace Holloway would much rather be asked than surprised.
2. If you use a personal account on a general chat assistant, look for its data or history settings and read how your conversations may be used. Chapter 3 explains what to look for. Terms change, so check the current version.
3. Choose a topic you know very well. This is essential: you will be judging the output's accuracy, so pick something where you would instantly spot an error.

Step 1: Try a vague prompt (2 minutes)

Start by asking the way many people first do — briefly and without context.

Weak prompt

```
Give me a lesson on photosynthesis.
```

Read the result. You will probably get something that is broadly reasonable but generic: an outline that could be for any age group, with no sense of how long the lesson is, what the learners already know, or what they need to be able to do by the end.

Step 2: Try a specific prompt (4 minutes)

Now give the same assistant the kind of information you would give a capable colleague who was covering your class. Here is how Lena might write it.

Effective prompt

```
You are an experienced secondary science teacher.

I teach a Year 9 class (ages 13-14) of 28 students. Three are English-language
learners and two have dyslexia. They already know that plants need light and
water, but most have not yet learned the word equation for photosynthesis.

Suggest a warm-up activity (about 8 minutes) that introduces photosynthesis and
finds out what students already believe about how plants get their food.

Requirements:
- Use simple, clear language suitable for the whole class.
- Include one visual or hands-on element.
- List two common misconceptions I should listen for.
- Keep your answer under 250 words.
```

Compare the two results side by side. The second should be more focused, closer to usable, and more obviously designed for a particular group. The difference comes entirely from the context you provided — none of which identified a real student.

Step 3: Check and refine (4 minutes)

Now act as the professional reviewer. Check the science. A correct simple explanation is that plants use light energy to convert carbon dioxide and water into glucose (a sugar) and oxygen, and that this happens mainly in the leaves, in structures called chloroplasts that contain chlorophyll. A common misconception worth listening for is that plants get their food from the soil. If the assistant said anything different, note it — you have just seen why verification matters.

Then ask one follow-up question to refine the result, for example:

Follow-up prompt

```
That is helpful. Please rewrite the student-facing instructions so that each
step is one short sentence, and add a small glossary of 4 key words with simple
definitions.
```

Here is the kind of short output you might see. Your result will differ, because AI tools rarely give the same answer twice.

Sample AI output

```
Student steps:
1. Look at the two plants on your table.
2. One plant had light. One plant was kept in the dark.
3. With your partner, write down what is different.
4. Answer: Where do you think a plant gets its food?

Glossary:
- Photosynthesis: how plants make their own food using light.
- Chlorophyll: the green substance in leaves that absorbs light.
- Glucose: a type of sugar the plant makes and uses for energy.
- Carbon dioxide: a gas in the air that plants take in.
```

Lena would check this before using it. The definitions are simplified but accurate for the level. She might decide the plants activity needs preparation a week in advance, or swap it for photographs — a decision only she can make, because only she knows her lab and her timetable.

A version for trainers

If you work in workplace training, try the same three steps with your own topic. Here is Jordan's pair.

Weak prompt

```
Write a training activity about customer service.
```

Effective prompt

```
You are an experienced trainer in healthcare customer service.

I run a half-day onboarding workshop for new front-desk staff at a regional
healthcare provider. Most have customer-facing experience in retail or
hospitality but none in healthcare.

Design a 15-minute opening activity that helps them see how speaking with a
worried patient differs from serving a retail customer.

Requirements:
- Suitable for 12 adult learners working in pairs.
- No real patient information; invent any details.
- Include 2 short scenarios and 3 debrief questions.
- Plain language, warm and respectful tone.
```

Jordan's review questions would be: Are the scenarios realistic for our front desk? Do any of them stray into clinical advice that front-desk staff should not give? Does the tone match our organization's values? Those questions are the professional judgment that makes the output safe to use.

> **TIP**
>
> Keep a simple note of your first try: the weak prompt, the effective prompt, what was good, and what you had to fix. This becomes the first entry in your personal prompt library, which we build properly in Chapter 22.

What We Learned from the First Try

This small exercise demonstrates most of the big ideas in this book. Context changes quality: the more clearly you describe learners, goals, and constraints, the better the draft. Iteration is normal: a short follow-up turned a reasonable answer into a more usable one. Verification is not optional: even a good-looking output needs a knowledgeable human check. Privacy can be protected: you described your class accurately without naming anyone. And judgment remains yours: the AI suggested; Lena and Jordan decided.

If your first try was disappointing, do not worry. Look at what you gave the tool. Was the request specific? Did you say who the learners are, what they already know, and what format you wanted? Chapter 4 is devoted to exactly this skill, and most people improve quickly once they see a few good patterns.

How Brightwater Is Approaching AI

It helps to see how different people in the same organization start from different places, because your own colleagues probably do too.

At Brightwater Secondary, Grace Holloway's first step was not to buy anything. She set up a small working group — including Lena as an enthusiastic early adopter, Sam as a constructive skeptic, and Mei to represent learning support and accessibility — to agree on principles before tools. Their early principles were simple: no personal student data in unapproved tools; AI outputs always reviewed by a teacher; students under 18 use AI only through approved, supervised routes and within each tool's own age requirements; and clear, honest communication with parents. Chapter 3 follows Grace's tool approval process, and Chapter 25 returns to her policy and professional development work.

At Brightwater Health Academy, Tom Becker started with a different question: "Which of our training tasks are slow, repetitive, and low risk?" Tom asked Jordan and Priya to keep a two-week log of their preparation time. Tasks such as drafting scenario variations, writing first-draft knowledge-check questions, and turning meeting notes into action lists came up again and again. Tasks such as confirming clinical accuracy or deciding who passes a competency check were marked as "human only." That simple log became the basis for a pilot, and Chapter 24 shows how Tom plans to evaluate it using Kirkpatrick's four levels of training evaluation.

Neither approach is the single right way. The common thread is that both leaders started with purpose and people, not with technology.