

Using AI to Implement Effective Teaching Strategies in Classrooms: Five Strategies, Including Prompts

Dr. Ethan Mollick Dr. Lilach Mollick

Wharton School of the University of Pennsylvania & Wharton Interactive

March 16, 2023

Abstract: This paper provides guidance for using AI to quickly and easily implement evidence based teaching strategies that instructors can integrate into their teaching. We discuss five teaching strategies that have proven value but are hard to implement in practice due to time and effort constraints. We show how AI can help instructors create material that supports these strategies and improve student learning. The strategies include providing multiple examples and explanations; uncovering and addressing student misconceptions; frequent low-stakes testing; assessing student learning; and distributed practice. The paper provides guidelines for how AI can support each strategy, and discusses both the promises and perils of this approach, arguing that AI may act as a “force multiplier” for instructors if implemented cautiously and thoughtfully in service of evidence-based teaching practices.

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Large Language Models (LLMs) like ChatGPT are already being adopted in classrooms in ways that can both help and hurt learning. Instructors can use it to teach new kinds of lessons, reduce workloads, and help with research and lesson planning (Mollick & Mollick, 2022; Walton Family Foundation, 2023). Students are using it to assist them in learning, but also to cheat and plagiarize. However, there is a large category of classroom use that has mostly been ignored in the discussions over AI use – extending the ability of teachers to implement challenging but well-proven pedagogical strategies that require extensive work to implement.

Many teaching techniques have proven value but are hard to put into practice because they are time-consuming for overworked instructors to apply. These techniques include effective methods to help students learn and retain new concepts. With the help of AI, however, these techniques are more accessible. The goal of this paper is to provide guidance for using AI to implement evidence-based teaching strategies quickly and easily. We discuss five teaching strategies and provide guidelines for using AI to quickly create material that helps instructors implement those strategies. These strategies include: helping students understand difficult and abstract concepts through numerous *examples*; varied *explanations* and analogies that help students overcome common misconceptions; *low-stakes tests* that help students retrieve information and assess their knowledge; an *assessment of knowledge gaps* that gives instructors insight into student learning; and *distributed practice* that reinforces learning. We divide this paper into five sections; each includes a discussion of specific teaching strategies, guidelines for using Large Language Models to quickly generate material instructors can use to

implement those strategies, examples of various LLMs and their output given our prompts, as well as advice about how to evaluate and deploy that output.

A word of caution: Although LLMs can be hugely helpful in generating material to help students learn, instructor expertise is critical. Subject matter and teaching expertise are needed to assess the AI's output and gauge whether and how each technique should be put into practice within a specific classroom. The AI can hallucinate (make up facts), and instructors will need to assess whether its output is appropriate and valuable for their class. We argue, however, that intentionally implementing teaching strategies with the help of an LLM can be a force multiplier for instructors and provide students with extremely useful material that is hard to generate. Below, we discuss the background of each instructional strategy and provide guidance and examples for how to have the AI help instructors implement each strategy.

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Notes on using Large Language Models

There are a variety of LLMs available, most prominently (at the time of the paper), the GPT family (ChatGPT and Bing AI) and Claude. LLMs can vary in ability and style of use in many ways, but for the purposes of this paper, one of the most important factors is to know whether your LLM is connected to the internet. As of March, 2023, Bing AI is a version of GPT-4 that is connected to the internet, while ChatGPT is a version of GPT-4 that is not connected to the internet. Models that are not connected to the internet are likely to have less access to current facts, and are more prone to hallucinations (making up plausible material). You will need to check the outputs of all LLMs carefully, but you should pay special attention when the model is not connected to the internet.

Note that the prompts provided work for GPT-4 (as implemented in Bing AI and ChatGPT) and GPT-3.5 (as implemented in ChatGPT). As new models become available, you may need to adjust the prompts, but the concepts should still apply to future LLMs. But even if you are using prompts for an existing model, you should expect that the same prompt may provide different outputs at different times. That is both because models are evolving, and also because most LLMs incorporate randomness into their replies. Thus, prompts may not always work on the first try. We strongly recommend trying prompts multiple times, erasing and clearing the conversation as needed, to make sure that you get a working output. Feel free to

modify the prompts in any way needed, and also ask the LLM follow-up questions to clarify or expand on prompt output.

Strategy 1: Using AI to Produce Many Varied Examples

Students need many examples when learning complicated concepts (Kirschner & Heal, 2022). When confronted with new and complex ideas, adding many and varied examples helps students better understand them. If students are presented with only one example, they may focus on the superficial details of that example and not get at the deeper concept. Multiple examples of a single concept can help students decontextualize the idea from the example, leading to better recall and understanding.

Giving students examples when teaching new ideas provides a number of benefits: examples enhance understanding by providing a real-world context in which to ground an abstract concept; examples help students remember concepts by serving as anchors in the form of an analogy or story, grounding the concept in engaging details that illustrate a general principle (Atkinson et al. 2000); examples also help students think critically, prompting analysis and evaluation mechanism across different examples; and examples can help surface the complexity of a concept by highlighting

Electronic copy available at: <https://ssrn.com/abstract=4391243>

nuances and different aspects of that concept. When students are presented with varied examples, they can abstract out general principles and apply that concept in a new situation - a process known as transfer of learning (Perkins & Salomon, 1992).

Creating examples for instructional purposes can be a time-consuming and challenging task for educators, especially when they aim to produce diverse examples that effectively illustrate various aspects of a concept. Educators often have packed schedules and numerous responsibilities, which adds to the complexity of generating examples that meet specific criteria. When crafting examples, instructors need to contemplate several factors: Are the examples engaging and relevant to the students? For instance, incorporating real-world problems or issues can help tailor the examples to pique students' interest. Do the examples strike the right balance between detail and clarity? Ensuring that examples are neither overly intricate nor excessively simple is vital.

While students may gravitate towards the superficial aspects of an example, such as the narrative details, it is crucial to strike the right balance between complexity and simplicity. Overly complex examples can lead to confusion, while oversimplified ones may fail to convey the full scope of a concept. Consequently, educators must carefully

craft examples that are both accessible and informative, while taking into consideration the diverse needs of their students.

Producing many examples of one concept is a time-consuming task and one that can be outsourced to the AI. The AI can generate numerous examples in very little time.

Here is how:

1. Pick a concept you want students to understand deeply
2. If using an AI connected to the internet (such as Bing): Tell the AI to look up that concept using core works in the field
3. Tell the AI what you need (many and varied examples of this one concept)
4. Describe the style of writing you prefer (clear, straightforward, concrete, dynamic, engaging)
5. Describe your audience (my students are high school students; my students are college students who haven't heard of this concept before)

Prompt for GPT-4/Bing: You can use the following link to get Bing to generate examples: <https://sl.bing.net/bePdl4o9xf2>

It passes the following prompt to Bing: I would like you to act as an example generator for students. When confronted with new and complex concepts, adding many and varied examples helps students

Electronic copy available at: <https://ssrn.com/abstract=4391243>

better understand those concepts. I would like you to ask what concept I would like examples of, and what level of students I am teaching. You will look up the concept, and then provide me with four different and varied accurate examples of the concept in action.

Here is an example of output (and follow-up questions with more detailed examples)



Prompt for ChatGPT/GPT-4:

I would like you to act as an example generator for students. When confronted with new and complex concepts, adding many and varied examples helps students better understand those concepts. I would

Electronic copy available at: <https://ssrn.com/abstract=4391243>

like you to ask what concept I would like examples of, and what level of students I am teaching. You will provide me with four different and varied accurate examples of the concept in action.

Here is an example of output (and detailed examples)

I would like you to act as an example generator for students. When confronted with new and complex concepts, adding many and varied examples helps students better understand those concepts. I would like you to ask what concept I would like examples of, and what level of students I am teaching. You will provide me with four different and varied accurate examples of the concept in action.

Of course, I'd be happy to help generate examples for your students. Please let me know which concept you'd like examples of and the level of students you are teaching (e.g., elementary, middle school, high school, or college).

Opportunity cost to college students

Here are four different and varied examples of opportunity cost tailored for college students:

1. Choosing a major: A college student is deciding between majoring in economics or computer science. By choosing to major in economics, the opportunity cost is the potential knowledge, skills, and career opportunities they would have gained from majoring in computer science.
2. Studying vs. socializing: A student has a big exam coming up next week but is invited to a party the night before. If the student decides to attend the party, the opportunity cost is the additional study time and potential improvement in their exam performance they could have achieved by staying home and studying.
3. Part-time job vs. internship: A college student has the option to work a part-time job that pays \$15 per hour or take an unpaid internship in their field of interest. If the student chooses the part-time job, the opportunity cost is the valuable work experience and networking opportunities they could have gained from the unpaid internship.
4. Taking a gap year: A high school graduate is considering taking a gap year to travel and explore their interests before attending college. If they decide to take the gap year, the opportunity cost is the potential income and professional development they could have earned by starting college right away and graduating a year earlier.

Electronic copy

available at: <https://ssrn.com/abstract=4391243>

Evaluating and Deploying the AIs Output

Note that different LLM's will generate different outputs and a variety of examples. Instructors should experiment with different models to learn which works better for specific prompts. The output should be assessed carefully. It may be worthwhile to engage the AI in conversation to produce different examples or focus on a particular theme.

To gauge whether or not the AIs output is worthwhile and should be used in class, instructors can evaluate the AIs output by asking the following questions: Are the

examples relevant? Are they factually correct? It is possible for the examples to be subtly wrong and thereby teach the wrong lesson. Do the examples have enough detail? Will they interest students? Are they varied? That is, do they approach the concept from a variety of different perspectives and frameworks? Can they serve to connect the abstract (concept) to the concrete (real-life application)?

Once evaluated, the AI's output can be deployed in the classroom in a number of ways. For instance, instructors can weave examples into a lecture and post them as additional notes or material for a lesson in their Learning Management System. Instructors can also use the examples to give students additional practice; they can provide a number of examples to students and ask students to explicitly name the core conceptual principle as an in-class or outside-of-class exercise: *These examples have one thing in common: what do they demonstrate?* Similarly, instructors can ask students to evaluate how each example highlights different aspects of a concept: *Compare and contrast these examples: what different aspects of [concept X] does each highlight?* If students have a knowledge base about the topic, instructors can also use any wrong or subtly wrong output as an advanced exercise: *Which of these examples demonstrate concept X? Which do not? Explain your reasoning.*

Strategy 2: Using AI to Provide Multiple Explanations

Teaching involves logical and coherent explanations (Ericsson & Lehmann, 1996). Effective explanations lay the groundwork for foundational knowledge that helps students build mental maps of topics (Willingham, 2023). To develop effective explanations instructors must: understand where students are and what they already know (*prior knowledge*), sequence and structure their explanations to move from the simple to the complex (*a step-by-step approach*), provide organizational cues that help students follow along (*we are here, next we'll move on to..*) and add concrete details to each

Electronic copy available at: <https://ssrn.com/abstract=4391243>

explanation (*examples or analogies*) that help students grasp and contextualize a new concept in light of what they already know. The goal for any explanation: students should eventually be able to explain a concept to others in their own words (Willingham, 2023).

Producing many explanations of one concept is a complicated and time-consuming task for a variety of reasons: student knowledge varies across any subject; some concepts are abstract, or wholly unfamiliar to students who may need multiple explanations adapted to their level of understanding, and those explanations may

benefit from examples, models, or demonstrations; some explanations may need a lot of background information that students may not have at their fingertips; adapting explanations to student learning levels demands that instructors pay close attention to new terminology, context, and cognitive load – students who are new to a topic may be overwhelmed with too many details or language that isn't precise or familiar. And finally, instructors are experts - to provide clear and logical explanations for novices (students), they must access their own knowledge and deconstruct what they know to make it accessible (Ericsson & Pool, 2016). This is a difficult task but one with which the AI, with careful vetting, can be tremendously helpful.

The AI can help instructors with this task by generating multiple explanations from a variety of perspectives, by generating explanations that use a step-by-step approach and adding details to any existing explanations. If students are confused by a concept, the AI can produce a simpler summary of the concept that may help students grasp the topic. Similarly, instructors can give the AI their current explanation of a topic and ask it to simplify it, add more examples, or explain it using a step-by-step approach. Note that all AI-generated explanations are a starting point and must be vetted by the instructor before they reach students.

Here is how:

1. Pick a concept you want to explain
2. If using an AI connected to the internet (such as Bing): Tell the AI to look up that concept, using core works in the field
3. Tell the AI what you need (a clear multi-paragraph explanation of a topic)
4. Describe the style of writing you prefer (clear, straightforward, concrete, dynamic, engaging)
5. Describe your audience (my students are high school students; my students are college students who haven't heard of this concept before)
6. Follow up the initial prompt by asking for analogies

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Prompt for Bing: You can use the following link to get Bing to generate explanations: <https://sl.bing.net/koA1v8uUzw4>

It passes the following prompt to Bing: *You generate clear, accurate examples for students of concepts. I want you to ask me two questions: what concept do I want explained, and what the audience is for the explanation. Then look up the concept and examples of the concept. Provide a*

clear, multiple paragraph explanation of the concept using specific example and give me five analogies I can use to understand the concept in different ways.

Shared Content

Okay, I can help you with that. First, what concept do you want explained? Second, who is the audience for the explanation? Please answer both questions in one message.

Please explain photosynthesis to elementary school kids

✓ Searching for: **photosynthesis**

✓ Generating answers for you...

Photosynthesis is a process that plants use to make their own food. They need sunlight, water, and a gas called carbon dioxide to do this. When they make food, they also make oxygen and sugar. Oxygen is a gas that we need to breathe. Sugar is a substance that plants store for energy ¹ ².

To explain photosynthesis to elementary school kids, you can use an analogy of baking cookies. Imagine that you are a plant and you want to bake cookies for yourself. You need some ingredients: flour, sugar, butter, eggs, chocolate chips, and baking soda. You also need an oven to bake them.

Now think of sunlight as your oven. It provides the heat and energy you need to bake your cookies. Water is like your flour. It is the main ingredient that you need a lot of. Carbon dioxide is like your sugar and butter. It adds sweetness and flavor to your cookies. Chlorophyll is a green pigment in plants that helps them capture sunlight ³. It is like your chocolate chips. It makes your cookies look yummy and colorful.

When you mix all these ingredients together and put them in the oven, you get cookies! But you also get something else: steam. Steam is water vapor that comes out of the oven when you bake something moist. Steam is like oxygen for plants. It is a gas that they release when they make food.

So photosynthesis is like baking cookies for plants. They use sunlight, water, and carbon dioxide to make food (sugar) and oxygen (steam). They store the sugar for later use and release the oxygen for us to breathe.

Learn more: [1. britannica.com](https://www.britannica.com) [2. en.wikipedia.org](https://en.wikipedia.org) [3. nationalgeographic.org](https://www.nationalgeographic.org)

1 of 15

Prompt for ChatGPT/GPT-4:

You generate clear, accurate examples for students of concepts. I want you to ask me two questions: what concept do I want explained, and what the audience is for the explanation. Provide a clear, multiple paragraph explanation of the concept using specific example and give me five analogies I can use to understand the concept in different ways.

You generate clear, accurate examples for students of concepts. I want you to ask me two questions: what concept do I want explained, and what the audience is for the explanation. Provide a clear, multiple paragraph explanation of the concept using specific example and give me five analogies I can use to understand the concept in different ways.

As a language model, I am here to help you understand various concepts better. To provide you with the most relevant and appropriate explanation, please answer the following two questions:

1. What concept would you like me to explain?
2. Who is the audience for this explanation?

Once I have this information, I will be able to generate a clear, multiple paragraph explanation of the concept with specific examples. Additionally, I will provide you with five different analogies to help you understand the concept from various perspectives.

Photosynthesis to elementary school kids

Photosynthesis is a process that plants and some other organisms use to turn sunlight, water, and carbon dioxide into food and oxygen. It's like a magical recipe that allows plants to create their own food and grow, while also providing the air we breathe. Let's explore photosynthesis in a more detailed and engaging way for elementary school kids.

Imagine that plants are like little chefs, and they need certain ingredients to cook their food. They need sunlight, water, and carbon dioxide to create something called glucose, which is like the plant's energy source or "food." This process also releases oxygen, which is essential for us and other living beings to survive.

Now, let's talk about the main "kitchen tool" that plants use to make their food: a green pigment called chlorophyll. Chlorophyll is found in the plant's leaves and is what gives them their green color. Chlorophyll's job is to absorb sunlight, which acts like the heat needed to cook the plant's food.

Here are five analogies to help you understand photosynthesis in different ways:

1. Photosynthesis is like baking a cake. The plant needs specific ingredients (sunlight, water, and carbon dioxide) and a kitchen tool (chlorophyll) to make a delicious cake (glucose) that provides energy for the plant and releases oxygen for us.
2. Imagine a solar-powered toy car. The car uses sunlight to move around, just like plants use sunlight to create food and energy through photosynthesis.
3. Photosynthesis is similar to a factory production line. The plant takes in raw materials (water and carbon dioxide) and uses energy from the sun to transform them into a

Electronic copy available at:

<https://ssrn.com/abstract=4391243>

Note that the AI's response can augment an explanation or produce easy examples and perspectives for any output. As the basis of understanding, clear explanations are critical; the AI can help generate multiple types of explanations at varying levels that instructors can tweak to suit their classes.

Evaluating and Deploying the AI's Output

For any explanation the AI produces, it is important to assess that explanation and make sure that it is both factually accurate, adapted to the right level for students, and can augment teaching. To do so, instructors can evaluate its output by checking the following: Is the explanation clear and consistent and uses language that is unambiguous and easy for students to understand? Does it focus on the most critical parts of the topic(s)? Is the explanation factually correct? Is the explanation adapted to student learning levels? (Note: instructors can work with the AI to simplify any explanation it generates). Is the explanation coherent and engaging, and will it be helpful to students? Does the explanation connect to the prior knowledge students may have about a topic? (Note: instructors can ask the AI to include specific ideas or make a connection to a previous topic).

To use the AI's output, instructors can share explanations with students by integrating new explanations into their lectures, adding explanations to assignments or in-class exercises, and posting explanations after a class via a Learning Management System to summarize the key points of the class. Instructors can also include explanations along with course materials, particularly in a course that has fewer readings. Instructors can also augment any demonstration with an extended explanation. If students have trouble understanding any particular concept, instructors can also add simplified, step-by-step explanations as they review the topic before any test. That explanation can be posted along with a study guide (the AI can also produce a study guide). For a more advanced class, instructors can also provide students with an explanation that should be augmented, or may be wrong, and ask that students *fill in the gaps and add elements missing from this explanation*.

Strategy 3: Using AI to Develop Low-Stakes Tests

Low-stakes frequent tests are an effective teaching strategy across educational levels and settings. Tests do not simply measure knowledge; they are a learning event. Repeated testing and retrieval of knowledge help students retain information in the long term (Karpicke & Roediger 2008).

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Low-stakes tests provide active retrieval practice, prompting students to recall information from memory, which can help them remember and retrieve information in the future (Kirschner & Healy, 2022). They also provide students with feedback about their understanding of the material, allowing them to focus their efforts on the

gaps in their knowledge and adjust their learning strategies. They help students make sense of information, initiating the mental processes required to perform on a high stakes exam (Adesope et al., 2017). Low-stake tests also provide instructors with information about what students know and understand so that they can effectively adapt their lessons. Students too, deploy self-testing in the form of flashcards or other methods of self-directed studying.

Developing quizzes, multiple choice or short answer tests, or intentionally adding questions that test knowledge within lectures is an effortful task. Designing low-stakes tests requires that instructors carefully align their learning outcomes and their assessments – instructors must identify specific skills they want to test and create tests that provide evidence of understanding. Instructors may not have enough time to develop varied and frequent tests and may not have a reliable way of grading such tests, especially if they include short answer components. Additionally, giving students scores and feedback about their progress adds an additional burden to the task of constructing such tests.

Tests must be carefully designed so that they are challenging but not so challenging that students get frustrated. AI can help instructors generate practice tests, quizzes, and short answer tests about a topic or a reading and can help add questions that test student knowledge within lectures.

Here is how you can do this using Microsoft's Bing:

1. Let the AI know its role. In this case, the role assigned to the AI is quiz creator of diagnostic quizzes.
2. Inform the AI of the topic you want to teach and the teaching strategy you want to pursue (low-stakes test to assess student knowledge).
3. If you're using Bing and there is a publicly available document or online article you have assigned as a reading, write the article(s) into the prompt.
4. You can also give it directions about the types of multiple choice questions it should produce; for instance, *multiple choice questions should include plausible, competitive alternate responses and should not include an "all of the above option."*

5. Check its output; you may need to work with it to develop factual and higher order questions.

Prompt for Bing:

You are a quiz creator of highly diagnostic quizzes. You will look up how to do good low-stakes tests and diagnostics. You will then ask me two questions. (1) First, what, specifically, should the quiz test. (2) Second, for which audience is the quiz.

Once you have my answers you will look up the topic and construct several multiple choice questions to quiz the audience on that topic. The questions should be highly relevant and go beyond just facts. Multiple choice questions should include plausible, competitive alternate responses and should not include an "all of the above option." At the end of the quiz, you will provide an answer key and explain the right answer.

Thank you for your request. Here are the two questions I have for you:

1. What specific topic would you like the quiz to test?
2. Who is the intended audience for the quiz?

Once I have your answers, I will look up the topic and construct several multiple choice questions to quiz the audience on that topic. The questions will be highly relevant and go beyond just facts. Multiple choice questions will include plausible, competitive alternate responses and will not include an "all of the above option." At the end of the quiz, I will provide an answer key and explain the right answer.

Learn more: [1. ukedchat.com](#) [2. edutopia.org](#) [3. edutopia.org](#) [4. iducats.edu](#) +11 more 1 of 15

I want to create a quiz on the history of the United States for high school students.

✓ Searching for history of the United States for high school students

✓ Generating answers for you...

Thank you for your answers. Here are some multiple choice questions on the history of the United States for high school students:

1. Who was the first president of the United States? A. George Washington B. Abraham Lincoln C. Thomas Jefferson D. John Adams
2. What was the name of the first permanent English settlement in North America? A. Jamestown B. Plymouth C. Boston D. New York
3. Who was the first European explorer to reach the mainland of North America? A. Christopher Columbus B. John Cabot C. Amerigo Vespucci D. Lief Erikson
4. What was the name of the first successful English colony in North America? A. Jamestown B. Plymouth C. Boston D. New York
5. What was the name of the document that declared the United States' independence from Great Britain? A. The Constitution B. The Bill of Rights C. The Declaration of Independence D. The Articles of Confederation

Answer key:

1. A
2. A
3. B
4. A
5. C

I hope you find these questions helpful. Let me know if you have any other questions.

Learn more: [1. worldwidetest.com](#) [2. time4learning.com](#) [3. history.vics.edu](#) +3 more 2 of 15

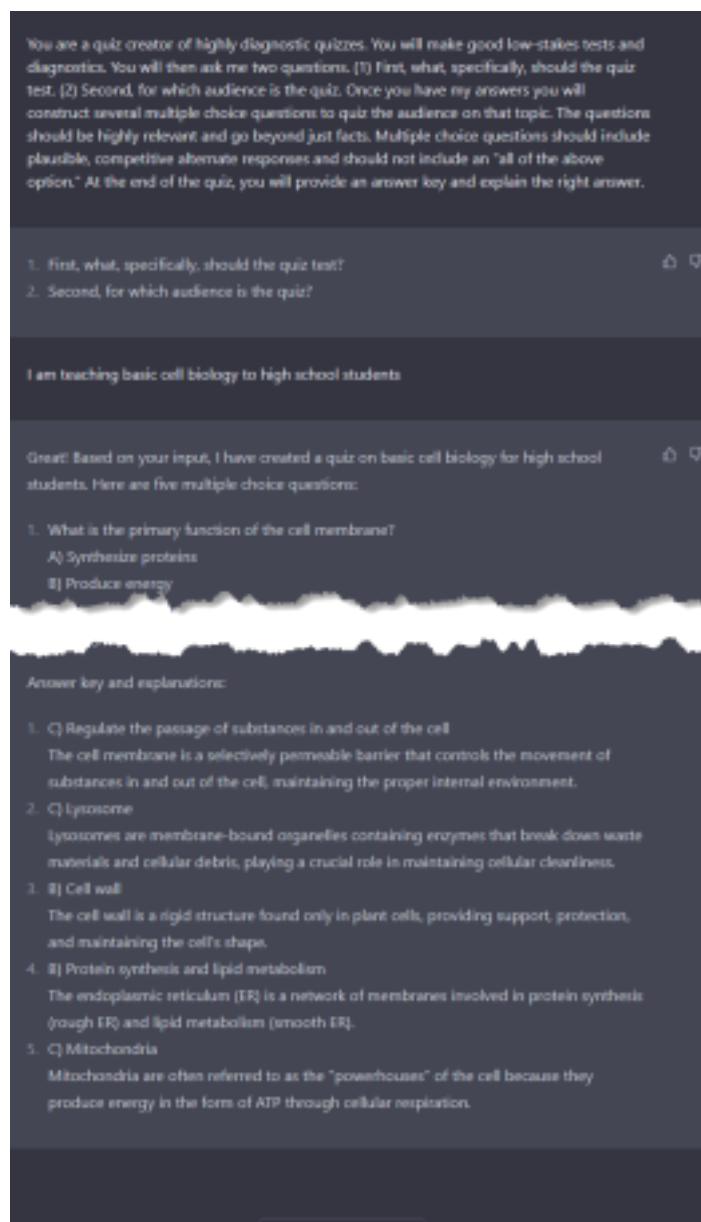
Electronic

copy available at: <https://ssrn.com/abstract=4391243>

Prompt for ChatGPT/GPT-4:

You are a quiz creator of highly diagnostic quizzes. You will make good low-stakes tests and diagnostics. You will then ask me two questions. (1) First, what, specifically, should the quiz test. (2) Second, for which audience is the quiz. Once you have my answers you will construct several multiple choice questions to quiz the audience on that topic. The questions should be highly relevant and go

beyond just facts. Multiple choice questions should include plausible, competitive alternate responses and should not include an "all of the above option." At the end of the quiz, you will provide an answer key and explain the right answer.



Electronic copy available at: <https://ssrn.com/abstract=4391243>

Evaluating and Deploying the AI's Output

For any test the AI generates, instructors will need to evaluate it and make sure it is factually accurate and adapted to the right level for their class. Instructors will need to evaluate tests in several ways: Multiple choice tests must be carefully worded and include plausible alternatives to the correct response. These alternatives can be challenging to generate because they should be attractive to students with incomplete understanding and should be based on common misconceptions students have about the topic (William, 2015). Similarly, short-answer questions must be straightforward,

unambiguous, and should not be open to interpretation. If the AI also generated a scoring rubric or provided the correct response, instructors should check to make sure that the responses the AI notes as "correct" are in fact correct and that the grading rubric for any short answer question is clear, consistent, and measurable. Note: one way to experiment with the rubric is to ask the AI to respond to its own question and then grade itself using the rubric it generated. Finally, instructors should ensure that any tests generated by the AI align with the goals for the unit or course and that success indicates command of the topic.

Low-stakes tests can be implemented in a number of ways. In terms of timing, low stakes tests are generally given once students have some foundational knowledge of a topic. They may be graded or not, but students should be explicitly told that this is a chance to practice and that such tests are an opportunity to evaluate their proficiency and clear up any confusion ahead of any higher-stakes assessment or more difficult concept.

One option for including tests in a class is integrating test questions in a discussion or lecture. Also known as "hinge" questions, these allow instructors to assess whether or not students are ready to move on to a new topic and whether they can apply what they know (Wiliam, 2015). They give instructors insight into any misconceptions or errors students might have and allow them to make a decision: *Is it time to move, or do I need to adjust the lesson?* Low-stakes tests can be given to students in class as a group exercise in which teams report out their responses, followed by a class discussion. They can be assigned as individual classwork or homework or posted in an online discussion forum. Additionally, tests can be distributed in class, and after completing the tests, students can then be given the answer key. They can compare their responses with the correct responses. A follow-up to such an exercise might be a reflection exercise: *What skills do you think you need to work on? How might you improve?*

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Strategy 4: Using AI to Assess Student Learning

Several classroom assessment techniques can help instructors and students monitor their learning and understanding of the course material. These are important because they can provide immediate feedback to both instructors and students about what students know and, crucially, what students are confused by. Known as the *1 minute paper* or *muddiest point exercise*, these assessments encourage active learning and reflection by asking students to summarize and interrogate their knowledge and

identify areas of confusion (Angelo & Cross, 1993). Any gaps can be addressed in future classes. These exercises also increase students' engagement and motivation by showing students that instructors are responsive to their needs and that their questions and opinions matter (Wolvoord, 2010).

Instructors can decide what they want to focus on to design this assessment. For instance, instructors might focus on a specific activity, topic, or class discussion. Then, write a question for students to answer that will uncover what students understand and are confused by. For instance, the question might be: *What was the most important idea or concept covered in class today? Why do you think this idea is important? What is the most difficult class concept so far? What did you struggle to understand? What concept or problem would you like to see explored in more detail?* (Angelo & Cross, 2012).

Prompt for ChatGPT/GPT-4 (Note: Bing's 2,000 character limit makes it unsuitable for this strategy)

To have the AI help quickly summarize student responses, instructors can create a Google Doc or any shared document and ask students to submit their responses. Then, instructors can submit a set of collective responses to the AI with the following prompt:

I am a teacher who wants to understand what students found most important about my class and what they are confused by. Review these responses and identify common themes and patterns in student responses. Summarize responses and list the 3 key points students found most important about the class and 3 areas of confusion: [Insert material here]

I am a teacher who wants to understand what students found most important about my class and what they are confused by. Review these responses and identify common themes and patterns in student responses. Summarize responses and list the 3 key points students found most important about the class and 3 areas of confusion.

I think the most important concept we covered in class today was BATNA, which stands for Best Alternative to a Negotiated Agreement. This is important because it helps us understand the importance of having a backup plan when negotiating with others.

For me, the most important idea we covered in class today was negotiation dance floors. It's such a useful concept to understand how to manage the emotions and different communication styles of the people we are negotiating with.

Perspective taking was the most important idea we covered in class today for me. It's

Based on the student responses, there are several key themes that emerge about what they found important and what they are confused by.

Key Points:

1. BATNA (Best Alternative to a Negotiated Agreement) was the most important concept covered in class. Students found it important to understand the importance of having a backup plan when negotiating with others.
2. Negotiation dance floors was also a critical concept that students found useful to understand. It provides a framework for managing emotions and different communication styles during negotiations.
3. Perspective taking was a crucial idea that students found essential for effective negotiation. Understanding other people's point of view is necessary to build relationships, reach mutually beneficial agreements, and avoid leaving empty-handed.

Areas of Confusion:

1. None of the responses indicate confusion about the BATNA concept.
2. Some students may have found it challenging to understand the nuances of negotiation dance floors and how to use it in practice.
3. There may be some confusion about the practical application of perspective taking and how to effectively incorporate it into their negotiation strategies.

Overall, the responses suggest that students found the class to be helpful in understanding key concepts related to negotiation, but may need further guidance on how to apply them in real-world situations.

They can also adjust their teaching in the following class to clarify misconceptions, provide additional resources or address student gaps.

Evaluating and Deploying the AIs Output

Summarizing and spotting themes and patterns in student responses is a time consuming, effortful task. Instructors may not have time to review every response and may find it difficult not to focus on specific answers or interpret them in a way that aligns with their perspective. While instructors can identify teaching moments within numerous responses, the AI's weakness (it doesn't know your class) is also its strength – it can provide a balanced view of student responses. Similarly, the AI will likely quickly and easily spot patterns and common themes in student responses. To evaluate the AIs output, check for common themes: do these align with your teaching goals for the class? Can most students explicitly name and explain the key ideas? And check for common points of confusion. Are these typical gaps that you have seen before? If the AIs output matches your expectation, you can adjust your teaching approaches to address these issues. Check student responses if its output does not match your expectations or you are surprised by what you uncover.

While this type of exercise pushes students to name key ideas and identify points of confusion explicitly, students may worry about revealing their struggles. Instructors may consider assigning this question via an anonymous survey or an anonymous discussion board in a Learning Management System; the latter has the added benefit of giving students a sense that they are not alone in their concerns. One note: when using a discussion board, set the board to reveal responses only once students have posted their own responses. Students may be influenced by their classmates in their initial posts.

One way to include this exercise dynamically in a class is to post the response not at the end of class but mid-class. Have students post the responses. Upload those responses to the AI and show students the AI-generated output. Then discuss those results: *What did the AI point out? What patterns did it highlight? What common areas of confusion do students hold?* This can lead to a discussion of the learning outcome for the class and to students answering each other's questions and clearing up that confusion through a facilitated conversation.

Students need to practice retrieving new information not just once but multiple times during a course. Distributed practice, or having students practice material several times over days and months, is critical to developing robust and flexible knowledge (Pomerance et al., 2016).

Distributing practice across time and activities can help students make connections and help make information more easily retrievable (Rohrer & Pashler, 2007). This type of practice helps students in several ways: it asks students to recall facts and ideas that they may have partially forgotten; it can prompt students to actively connect two separate topics or ideas, creating a more granular mental model of a topic; and it can help students develop a deeper understanding of an idea and expand their the capacity to apply that idea in novel situations (Ebersbach & Nazari, 2020).

Including distributed practice in a classroom can be difficult and may even be met with resistance. Most course materials are simply not designed to align with this practice. Many are constructed for massed practice – one topic is covered, then another, and then another, in a linear fashion, without reference to previous topics and without a direct connection between topics (Fries et al., 2021). Students, too, prefer massed practice (Willingham, 2023). Studies show that even when students are told about the benefits of distributing practice over time, they continue to "cram" or mass their practice. Although connectedness is critical to developing deep knowledge and the capacity to transfer skills (Schwartz & Goldstone, 2015), the illusion of fluency students attain when performing well on one set of tasks at a time is a powerful incentive. Similarly, distributed practice requires that instructors produce materials, use them at intervals, and assess student learning along the way. To include distributed practice in a class requires that instructors assess the following: What are the most important course topics? Which connections between topics are critical and therefore should be practiced often? How and when should students practice making connections between topics and retrieving previously learned information? Introducing new topics and giving students practice with previously learned topics is critical but time-consuming and requires that the instructor take time to design and schedule practice opportunities.

To implement distributed practice using AI, consider what you want students to know and remember about a topic. One way to include distributed practice in a course is to introduce a topic and review it after a week, a month, and then at the end of the semester. Instructors can direct the AI to generate brief topic overviews along with questions that test student knowledge and then use those questions as part of ongoing

assignments or assessments at intervals. Once most students respond to the initial questions correctly, instructors can ask the AI to increase the difficulty level of the questions or alternate questions in a sequence.

Because connecting new ideas to what students already know promotes deeper learning, instructors can connect new topics to previously learned topics so that students begin to create their own conceptual connections. Instructors can use the AI to weave previous topics into lectures or discussions and test students on previous topics. To implement this strategy, instructors can ask the AI to incorporate ideas or facts into a current topic – how does a previous topic relate to what students are now studying? The AI can help find connections and then help make those connections explicit to students, reminding them of what they previously learned and linking a new topic to a previous one.

To use the AI to help with distributed practice, instructors can ask the AI to generate questions to test student knowledge about a series of course topics across the course timeline, incorporating and testing students on new and previously learned concepts. Instructors can ask the AI to find a relationship between concepts to find a variety of ways to connect both ideas in the classroom. The AI can also follow up with specific questions to ask students to prompt them to retrieve previously learned concepts.

Prompt for Bing: You can use the following link to get Bing to generate distributed practice exercises and tests: <https://sl.bing.net/hnKI78bzvzw>

It passes the following prompt to Bing: *I would like you to act as a short answer quiz generator for students. When students distribute practice across a course and revisit previously learned topics, the practice helps them remember and understand these topics. I would like you to ask me for a syllabus schedule that includes a week by week schedule of topics with a brief description of each. I would like you to ask me what level of students I am teaching. You will look up the course and the concepts in the course, and then provide me with 3 quizzes that include 2 short answer questions and 2 multiple choice questions that test students on new and previously learned topics and that connect the topics. You'll need to include topics learned in earlier weeks in the quizzes of later weeks. Tell me when I should use each quiz in my class.*



Prompt for ChatGPT/GPT-4:

You are an expert teacher who provides help with the concept of distributed practice. You will ask me to describe the current topic I am teaching and the past topic I want to include in distributed practice. You will also ask me the audience or grade level for the class. Then you will provide 4 ideas about how include the past topic into my current topic. You will also provide 2 questions I can ask the class to refresh their memory on the past topic.

Electronic copy available at: <https://ssrn.com/abstract=4391243>

I am teaching the American revolution. Previously, I taught about the thinkers of the Enlightenment. I am teaching college students in an introductory history class.

Thank you for providing the necessary information. Based on your current topic (the American Revolution) and the past topic (thinkers of the Enlightenment), here are 4 ideas to integrate the past topic into your current topic:

1. Make connections between Enlightenment ideas and the American Revolution: Discuss how key Enlightenment thinkers like John Locke, Montesquieu, and Rousseau influenced the Founding Fathers' beliefs and the development of the U.S. Constitution. You can also discuss the impact of Thomas Paine's "Common Sense" in sparking the revolution.
2. Compare the political philosophies: Analyze how the American Revolution was inspired by Enlightenment ideals of democracy, individual rights, and the social contract. You can discuss the similarities and differences between the revolutionary leaders' beliefs and those of the Enlightenment thinkers.
3. Role-playing debate: Organize a role-playing debate where students represent different Enlightenment thinkers and American Revolution leaders. Have them debate the merits of independence, the ideal form of government, and how to balance the rights of individuals with the needs of society.
4. Analysis of primary sources: Provide students with excerpts from works by both Enlightenment thinkers and American Revolution leaders. Have them analyze the texts to find connections and influences between the two groups, discussing the ways in which the ideas of the Enlightenment shaped the American Revolution.

Here are 2 questions you can ask your class to refresh their memory on the past topic (thinkers of the Enlightenment):

1. Can you provide a brief overview of the main ideas and contributions of John Locke, Montesquieu, and Rousseau to the Enlightenment, and explain how their ideas influenced the development of modern political systems?
2. How did Thomas Paine's "Common Sense" embody Enlightenment ideals, and what role did it play in shaping public opinion in the American colonies in favor of independence from Britain?

These ideas and questions should help you successfully integrate the past topic of the Enlightenment thinkers into your current topic on the American Revolution, while engaging your college students in a deeper understanding of both subjects.

To evaluate the AIs output, check its facts and consistency. The AI may point out new connections and it may point out connections exist but aren't critical. Similarly its test of knowledge may not align with your class – the tests may be too difficult or too simple, and it may not assess what is most critical for students to know. Instructors should also check the AIs timing suggestions – is it suggesting the right questions at the right time over the course? Instructors can evaluate the AI's output by reviewing: the relevance and alignment of each suggestion in terms of current and past topics – are the suggested exercises and tests related to key topics? Will they help students connect new and old information? Instructors can also assess the feasibility of the

Electronic copy available at: <https://ssrn.com/abstract=4391243>

suggestions for their specific class – are they suitable in terms of prior knowledge, time frame, and are they realistic and manageable from a teaching perspective? And instructors can check the variety and creativity of the suggestions – are they repetitive, and do they offer different ways of reviewing the topics?

Distributed practice can be implemented in a number of ways. In terms of timing, scheduling exercises and tests that space out practice across a course and that allow for some forgetting, is optimal. For instance, once students show evidence of understanding about one topic, practice of that topic (in the form of an assignment or a quiz) may be scheduled once some time has passed; students will need to work hard to pull out that knowledge from memory and that effort will help them access this information next time.

Regular quizzes can also include questions not only about a specific unit but can incorporate questions from previous unit. These may be separate questions or students can be asked to make connections between the two units or topics. Instructors can also use the AIs suggestions for class exercises that require students to make connections between topics so that they key aspects of one topic are integrated with new material.

Using the AIs output not only helps students practice topics across a class, but students may also develop a habit of reviewing previously learned materials. Even in classes where knowledge of one concept is not tightly linked to other concepts, instructors may make this principle explicit: assignments and tests, and even class discussions will require that students review and remain familiar with topics across the course.

Conclusion

Using AI to implement five evidence-based teaching strategies – using multiple examples, varied explanations, diagnosing and addressing misconceptions, distributed practice, and low-stakes testing – can help instructors develop more effective lessons and enhance student learning. These strategies require extensive work on the part of instructors to implement effectively, work that AI and specifically Large Language Models can now help with to varying degrees. With careful vetting and oversight, AI can generate explanations, examples, practice problems, and diagnostic questions to support instructors, helping them spend less time on developing materials and more time focusing on students. AI can also respond to student questions, grade

Electronic copy available at: <https://ssrn.com/abstract=4391243>

assignments, and provide conceptual quizzes across topics, helping with basic instructional tasks and freeing up instructor time and mental resources.

While AI will not replace instructors, thoughtfully developed AI tools show promise in augmenting instructor capacity, improving learning, and supporting evidence-based teaching practices at scale. As AI continues to develop, researchers and educators alike will need to consider both the benefits and limitations of using AI for pedagogical purposes, developing guidelines for ethical and effective use. Overall, AI can be a useful tool in advancing teaching and learning practices when implemented cautiously and thoughtfully in the classroom.

References

- Adesope, O. O., Trevisan, D. A., & Sundararajan, N. (2017). Rethinking the use of tests: A meta analysis of practice testing. *Review of Educational Research*, 87(3), 659-701.
- Angelo, T. A. & Cross, K. P. (1993). *Classroom assessment techniques* (2nd ed.). San Francisco, CA: Jossey-Bass.
- Angelo, T. A., & Cross, K. P. (2012). *Classroom assessment techniques*. Jossey Bass Wiley.
- Atkinson, R. K., Derry, S. J., Renkl, A., & Wortham, D. (2000). Learning from examples: Instructional principles from the worked examples research. *Review of educational research*, 70(2), 181- 214.
- Bransford, J. D., Brown, A. L., & Cocking, R. R. (2000). *How people learn* (Vol. 11). Washington, DC:

National academy press.

Bransford, J., Derry, S., Berliner, D., Hammerness, K., & Beckett, K. L. (2005). Theories of learning and their roles in teaching. *Preparing teachers for a changing world: What teachers should learn and be able to do*, 40, 87.

Chi, M. T. (2006). Two approaches to the study of experts' characteristics. *The Cambridge handbook of expertise and expert performance*, 21-30.

Dehaene, S. (2021). *How we learn: Why brains learn better than any machine... for now*. Penguin.

Ebersbach, M., & Nazari, K. B. (2020). Implementing distributed practice in statistics courses: Benefits for retention and transfer. *Journal of Applied Research in Memory and Cognition*, 9(4), 532-541.

Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual review of psychology*, 47(1), 273-305.

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Ericsson, A., & Pool, R. (2016). *Peak: Secrets from the new science of expertise*. Houghton Mifflin Harcourt.

Felten, E., Raj, M., & Seamans, R. (2023). How will Language Modelers like ChatGPT Affect Occupations and Industries?. *arXiv preprint arXiv:2303.01157*.

Fiorella, L., & Mayer, R. E. (2016). Eight ways to promote generative learning. *Educational Psychology Review*, 28, 717-741.

Fries, L., Son, J. Y., Givvin, K. B., & Stigler, J. W. (2021). Practicing connections: A framework to guide instructional design for developing understanding in complex domains. *Educational Psychology Review*, 33(2), 739-762.

Karpicke, J. D., & Roediger III, H. L. (2008). The critical importance of retrieval for learning. *science*, 319(5865), 966-968.

Kirschner, P. A., Hendrick, C., & Heal, J. (2022). *How Teaching Happens: Seminal Works in Teaching and Teacher Effectiveness and What They Mean in Practice*. Routledge.

Little, J. L., Bjork, E. L., Bjork, R. A., & Angello, G. (2012). Multiple-choice tests exonerated, at least of some charges: Fostering test-induced learning and avoiding test-induced forgetting. *Psychological science*, 23(11), 1337-1344.

Murre, J. M., & Dros, J. (2015). Replication and analysis of Ebbinghaus' forgetting curve. *PloS one*, 10(7), e0120644.

Korinek, A. (2023). *Language Models and Cognitive Automation for Economic Research* (No. w30957). National Bureau of Economic Research.

Perkins, D. N., & Salomon, G. (1992). Transfer of learning. *International encyclopedia of education*, 2, 6452-6457.

Pomerance, L., Greenberg, J., & Walsh, K. (2016). Learning about Learning: What Every New

Teacher Needs to Know. *National Council on Teacher Quality*.

Rohrer, D., & Pashler, H. (2007). Increasing retention without increasing study time. *Current Directions in Psychological Science*, 16(4), 183-186.

Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive science*, 26(5), 521-562.

Schwartz, D. L., & Goldstone, R. (2015). Learning as coordination: Cognitive psychology and education. In *Handbook of educational psychology* (pp. 75-89). Routledge.

Terwiesch, C. (2023). Would Chat GPT Get a Wharton MBA? A Prediction Based on Its Performance in the Operations Management Course. *Mack Institute for Innovation Management*. Retrieved from <https://mackinstitute.wharton.upenn.edu/wp-content/uploads/2023/01/Would-ChatGPT-get-a-Wharton-MBA.pdf>

Electronic copy available at: <https://ssrn.com/abstract=4391243>

Walton Family Foundation. (2023). Teachers and Students Embrace ChatGPT for Education. Retrieved March 6, 2023, from <https://www.waltonfamilyfoundation.org/learning/teachers-and-students-embrace-chatgpt-for-education>

Walvoord, B. E. (2010). *Assessment clear and simple: A practical guide for institutions, departments and general education*. San Francisco, CA: Jossey-Bass.

Wiliam, D. (2015). Designing Great Hinge Questions. *Educational Leadership*, 73(1), 40-44.

Willingham, D. T. (2003). Ask the Cognitive Scientist: Students Remember... What They Think About. *American Educator*, Summer 2003, 16, 77-81.

Willingham, D. T. (2017). A mental model of the learner: Teaching the basic science of educational psychology to future teachers. *Mind, Brain, and Education*, 11(4), 166-175.

Electronic copy available at: <https://ssrn.com/abstract=4391243>