

AI: What you need to know for your classroom

Physics Education Innovation Forum
Tue Mar 10, 2026 1-2pm

Janet Staker Woerner, Ph.D
John Martin, Ph.D

Today's session

1. Describe Generative AI (specifically LLMs)
2. Generative AI at UW-Madison (guidelines & policies)
3. Strategies (for yourself and students)
4. Considerations for your courses
5. Q & A

What is Generative AI

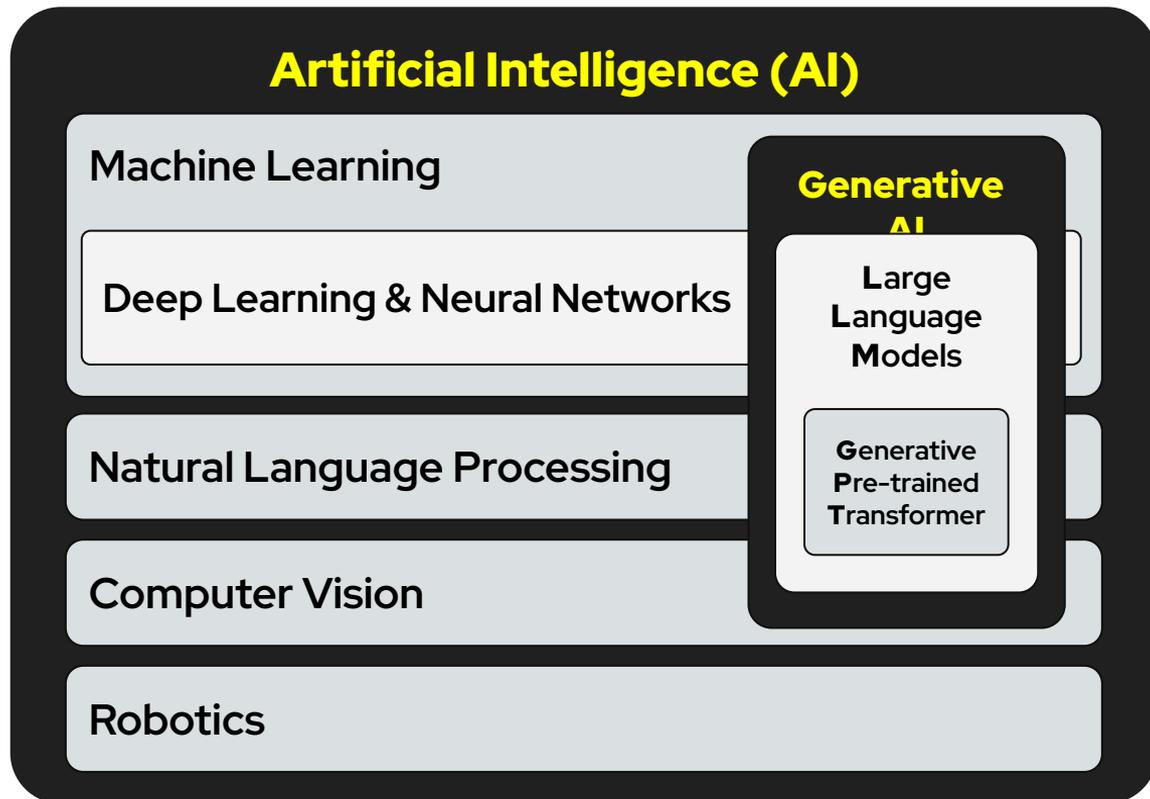
When we say “AI” what do we mean?

Artificial Intelligence (AI)

Coined in the 1950s for programs mimicking the behaviour of intelligent biological systems.

Generative AI

Pre-trained, Large Language Models (LLMs) that generate content using tokenization to predict what comes next (words, images, code, etc.).



Generative AI at UW

CTLM's Approach

Our students need knowledge, skills and experiences to help them thrive in a world powered by generative AI

Generative AI is at our fingertips and rapidly maturing. Keeping pace with the technical landscape, emerging capabilities, and new integrations will be a challenge but is essential, especially at a Research 1 university.

To participate in a complex and dynamic society that will increasingly depend upon AI, UW–Madison students will need skills such as prompt engineering, problem-solving, bias detection and intellectual curiosity. At the same time, they will need emotional intelligence, flexibility, and the ability to collaborate with humans and machines.

CTLM recognizes that there are a wide range of views about the use of generative AI in teaching, as well as about the broader social, environmental, and economic impacts of AI. We do not take a position on whether instructors should use AI. Rather, we support instructors in exploration, understanding, and thoughtful use or non-use of AI in ways that best serve their individual teaching and learning contexts.

Guiding Principles for Teaching

This guidance is based on the work of cross-disciplinary teams at UW–Madison including faculty, instructional academic staff, academic technologists and student affairs specialists. Campus guidance and resources will continue to evolve to reflect developments in this rapidly changing space.

[Expand all](#) | [Collapse all](#)

- + Advance accessibility and equity
- + Protect data and intellectual property
- + Consider educational uses
- + Adapt learning experiences and assessments
- + Discuss course expectations and academic integrity with students
- + Address potential misconduct using established policies and procedures

Campus-approved AI tools

Access generative AI tools at UW–Madison

UW–Madison has vetted and secured contracts for the generative AI services below. These tools are available university-wide for free and provide higher data security and privacy protection than public services. Please consider these options before exploring unvetted generative AI services for university work.



Google Gemini

Google Gemini is an AI chatbot designed to enhance productivity and creativity through natural language interactions. It can summarize and write text, translate languages, answer questions, generate images and write code. When you log in with your NetID, Gemini offers enterprise data protection, meaning it won't use your prompt data to train its large language models.

OK to use with: Public and internal data only

[Learn more about Google Gemini](#)



Microsoft 365 Copilot Chat

Copilot Chat is an AI-powered digital assistant that can help you with tasks like answering questions, creative writing, sample coding and creating images. When you log in with your NetID, Copilot Chat offers commercial data protection, meaning it won't use your prompt data to train its large language models.

OK to use with: Public and internal data only

[Learn more about Copilot Chat](#)



NotebookLM

NotebookLM is an AI-powered notebook designed to help you understand and work with your uploaded documents and other digital content. Ask questions about the source materials in your “notebook” to get summaries with citations and generate different types of content like study guides and audio overviews. NotebookLM is included with your UW–Madison Google account.

Okay to use with: public and internal data only

[Learn more about NotebookLM](#)



Zoom AI Companion

Zoom AI Companion can automatically create meeting summaries, identify action items and quickly answer questions about what happened during meetings. UW–Madison has 2 tiers of Zoom accounts—standard for most students and staff and elevated security for employees working with HIPAA-protected health data.

OK to use with: Public, internal and sensitive data (know your unit's guidance before using); Secure Zoom is also OK to use with restricted data, including HIPAA health data



Webex AI Assistant

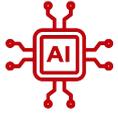
Webex AI Assistant is an AI-powered meeting tool that offers real-time translation, transcription, voice commands and post-meeting summaries, allowing you to stay on top of the conversation and make virtual meetings more efficient and inclusive.

OK to use with: Public, internal, sensitive and restricted data (know your unit's guidance before using)

Expectations for Students

Students are responsible for knowing their instructor's expectations when it comes to using AI tools. If it is unclear whether AI tools are allowed in a particular course or for an assignment, it is the student's responsibility to ask their instructor before using them.

Instructors' expectations will vary from course to course. If an instructor does not allow the use of AI tools and they suspect a student has used them, they will invite the student to meet with them about the possible violation of UWS 14.03(1)(b) Use of Unauthorized Materials.



Data Classification



Canvas *IgniteAI*

You may be hearing about new **Canvas** features released under the umbrella name *IgniteAI*.

IgniteAI is a new Canvas feature line from Instructure that groups together AI-driven functionality.

IgniteAI features are moving from the conceptual phase into the **Feature Preview stage**.

Feature Previews are optional for us to enable in our Canvas environment but are still under active development by the vendor. Therefore, we do not immediately enable them. In addition, several **IgniteAI** features are being offered under a “free for now, pay later” model.

What We're Hearing from Faculty

Assessment approaches that worked in the past may no longer be sufficient for the future.

Faculty are revisiting **fundamental questions** about what it means for students to learn.

Instructors need support to make meaningful, large-scale **changes** to their courses.

There is a growing need for more **discipline-specific examples** and use cases.

Faculty are looking for **clearer guidance** on when and how AI can be used.

There is strong interest in ongoing **conversation across disciplines**.

Teaching & Learning Strategies

Practice, Pain, and Potential

| Practice | Pain | AI helps | AI transforms |
|------------------------------------|--|--|--|
| Lecture & Peer Learning | One-way delivery; writing good conceptual questions is hard and time-consuming; misconceptions persist invisibly | Generate banks of clicker questions targeting <i>known</i> misconceptions with plausible distractors | – |
| Problem Sets | Students game them; formula-plugging; grading is brutal | Generate sets with varied surface features but identical deep structure (tests real understanding, not pattern matching) | Generate <i>elicit-confront-resolve</i> sequences – problems where misconceptions feel right at first, then produce obviously wrong results, then resolve into correct understanding. |
| Labs | Verification labs are pedagogically weak; designing good inquiry labs takes enormous time | Generate lab protocols with built-in uncertainty analysis scaffolding | Generate <i>fictional but physically consistent</i> datasets for scenarios – impossible before, now instant |
| Context-Rich Problems | Writing realistic scenarios is slow; hard to vary at scale | Generate topic-specific scenarios quickly | Generate <i>personalized</i> context variants for different student populations |
| Office Hours | Doesn't scale; TAs have uneven content knowledge | Prep students with AI-generated worked examples before office hours | AI Socratic tutor refuses to give answers and surfaces specific misconceptions driving errors |

Better Clicker Questions, Faster

- Writing good clicker questions is hard – the distractor has to be *wrong for the right reason* (not just wrong).
- Instructors know which misconceptions students stumble on.
- Generating a full question bank used to take hours. GenAI jumpstarts that process. (You still need to check, and refine based on specific challenges you witness for topics).

Prompt

I teach Physics 201. Students consistently confuse *net force* with *total force* – they'll add all forces including ones that cancel, then treat that sum as if it means something physical. Generate 5 clicker questions set on Mars that target this specific misconception. Each question should have one correct answer and three distractors – at least two distractors should be answers a student gets if they apply the misconception rather than random wrong answers. After each question include a one-sentence instructor note explaining which distractor maps to the misconception.

(Gemini) Clicker Questions Responses

These clicker questions are designed to isolate the "scalar-summation" misconception, where students treat the magnitude of forces as a cumulative "effort" rather than as vector quantities.

Question 1: The Stationary Rover

A 100 kg rover is at rest on a level Martian plain. The force of gravity pulls the rover down with 371 N, and the ground pushes up with a normal force of 371 N. What is the magnitude of the net force acting on the rover?

- A) 0 N
- B) 371 N
- C) 742 N
- D) 100 N

Instructor Note: Option C maps to the misconception that "net" means the arithmetic sum of all magnitudes ($\$371 + 371\$$), while Option B assumes the net force is simply equal to the magnitude of the active forces present.

Question 2: The Constant Velocity Descent A

lander is descending toward the Martian surface at a constant velocity of 5 m/s. The upward thrust from the engines is 2,000 N, and the downward gravitational force is 2,000 N. What is the net force?

- A) 0 N
- B) 2,000 N
- C) 4,000 N
- D) 5 N

Instructor Note: Option C reflects the scalar-summation error ($\$2,000 + 2,000\$$), whereas Option B targets students who believe the presence of motion implies a non-zero net force equal to the thrust.

Problem Sets: Patterns vs Contexts

- You (human) know where students struggle; AI does not. So you can create prompts that get to the right struggle for students.
- Students who "game" problem sets recognize *surface patterns* – same formula, different numbers. *Varied contexts* force them to identify the underlying physics.
- Generating dozens of structurally identical but contextually varied problems now takes one prompt.

Prompt

I teach introductory physics (Physics 201) where students consistently struggle with one specific thing: they identify forces by *what's moving* rather than *what's interacting*. They'll correctly find the gravitational force on a rover but miss the normal force from the ground entirely because "the ground isn't doing anything." Generate 5 problems testing Newton's 2nd Law ($F=ma$) set on Mars – rover navigation, habitat pressurization, astronaut EVA, dust storm equipment failure, robotic arm assembly – where correctly identifying *all* interacting forces is required to get the right answer. Students who only track moving objects will get wrong answers. Include numerical values and a worked solution that explicitly names each force pair.

(Gemini)

The decision you are making is whether to structure these physics problems to penalize "motion-based" force identification or to reward "interaction-based" force identification.

The fundamental misconception your students are facing is a failure to recognize **passive force** (constraints like normal force or tension) which only manifest to satisfy Newton's Second Law ($F_{\text{net}} = ma$) or to prevent interpenetration of matter.

The Socratic Question: If a student correctly identifies that an object is at rest ($a = 0$), but fails to identify the forces that cause that equilibrium, does their "correct" prediction of the motion actually indicate any understanding of the physics?

[\(See full response from Gemini\)](#)

Martian Physics Problem Set: Newton's 2nd Law

Environmental Constants for all problems:

- $g_{Mars} \approx 3.72 \text{ m/s}^2$
- ρ_{Mars} (Atmospheric Density) $\approx 0.020 \text{ kg/m}^3$ (surface)

1. Rover Navigation: The "Stuck" Wheel

Scenario: A 900 kg rover is attempting to climb a 15° basalt incline. The rover is moving at a constant velocity of 0.5 m/s . The student must calculate the required frictional force from the ground to maintain this motion.

- **The Trap:** Students often assume that because the velocity is constant, the "engine" is the only force, or they ignore the normal force because the ground is "stationary."
- **Worked Solution:**
 - **Force 1:** Weight ($W = mg$) acting straight down (3348 N).
 - **Force 2:** Normal Force (N) from the slope ($W \cos(15^\circ) = 3234 \text{ N}$).
 - **Force 3:** Static Friction (f_s) acting up the slope ($W \sin(15^\circ) = 866.5 \text{ N}$).
 - **Interaction Pair:** The rover pushes *down/back* on Mars; Mars pushes *up/forward* on the rover.
 - **Result:** Since $a = 0$, $F_{\text{net}} = 0$. f_s must exactly equal the downslope component of gravity.

Problem Sets: ECR (*elicit-confront-resolve* sequences)

- Effective ECR = knowing *what* students get wrong, AND *why* – the belief behind errors. Experienced instructors can capture this in plain language prompts.
- ECR works if the confrontation uses something students already trust as true. Designing them used to require physics education researchers.
- Students with misconceptions often feel confident at first – then get results that can't be right. That *productive confusion* is where learning happens.

Prompt

I've watched hundreds of students in Physics 201 work through Newton's 3rd Law. The misconception isn't just "heavier objects push harder" – students believe force is a property of an object, not an interaction between two objects. A Mars rover feels like it "has" more force than the rock it drives over. Generate a 6-problem ECR sequence set on Mars. Problems 1-2 should feel completely intuitive to a student who believes force is a property – they'll get right answers for wrong reasons. Problems 3-4 should feature a scenario where that belief produces numbers that violate something they know is true – something they can check. Problems 5-6 should reframe the rover-rock interaction explicitly as mutual, using the same scenario from problems 1-2 so students can see exactly where reasoning broke down. Include instructor notes naming the specific belief each problem is targeting.

(Copilot)

Here's a set of five Mars-based $F=ma$ problems where you *must* identify all interactions to get the right answer. I'll use $g_{\text{Mars}}=3.7 \text{ m/s}^2$.

[\(See full response from Copilot\)](#)

1. Rover navigation on a slope

A 250 kg Mars rover drives *up* a straight 15° rocky slope at constant speed. The rolling/kinetic friction coefficient between the rover's wheels and the ground is $\mu = 0.20$.

- Find the magnitude of the drive force the wheels must exert on the ground.
- List all forces on the rover and all Newton's 3rd-law force pairs involved.

Solution

Forces on the rover (object: rover):

- Gravity:** \vec{F}_g : Mars pulls rover down.
- Normal force:** \vec{N} : ground pushes rover perpendicular to slope.
- Friction:** \vec{f} : static/rolling friction from ground on rover, *up* the slope (opposes relative slipping).
- Drive force?** Careful: the rover's engine makes the *wheels* push on the ground; the ground pushes back on the wheels. On the rover as a whole, the *net* tangential contact force from the ground is \vec{f} up the slope. There is no separate "engine force" from nowhere.

Take axes: x along slope (up positive), y perpendicular to slope (out of surface).

- Normal force:**

$$\sum F_y = 0 \Rightarrow N - mg\cos\theta = 0 \Rightarrow N = mg\cos\theta$$

$$N = 250 \cdot 3.7 \cdot \cos 15^\circ \approx 250 \cdot 3.7 \cdot 0.966 \approx 894 \text{ N}$$

- Friction magnitude (limiting, assume enough to avoid slipping):**

$$f = \mu N = 0.20 \cdot 894 \approx 179 \text{ N}$$

(Gemini ECR)

The decision you are making is **how to architect a cognitive conflict sequence** that moves students from viewing force as an internal "impetus" or "power" to viewing it as a reciprocal interaction between two discrete entities.

In physics education research, this is the shift from **O-properties** (Object-based) to **R-properties** (Relational-based). Students often view a Mars rover as "having" force because it has a large battery and powerful motors.

One high-leverage question: **If a student correctly calculates the numerical value of a force using $F=ma$ but cannot identify the physical object providing that force, have they actually learned physics, or have they simply learned to "pattern match" variables to an equation?**

[\(See full response from Gemini\)](#)

The Martian Interaction Sequence (ECR)

Constants: $g_{Mars} = 3.71 \text{ m/s}^2$.

Part 1: The "Intuitive" Trap (Problems 1–2)

Goal: Allow students to succeed using "Force as a Property" logic to build a false sense of security.

Problem 1: The Rover's "Pushing Power"

A 1,000 kg rover is accelerating across a flat Martian plain at 2.0 m/s^2 .

- **Task:** Calculate the net force acting on the rover.
- **Result:** $F = 1,000 \cdot 2.0 = 2,000 \text{ N}$.
- **Instructor Note:** Students will correctly identify 2,000 N but will describe it as "the force the rover *has*" or "the force from the engine." They are attributing the force to the rover itself.

Problem 2: The Rock's "Resistance"

The rover hits a 50 kg rock, pushing it forward so the rock also accelerates at 2.0 m/s^2 (assume no friction for this moment).

- **Task:** Calculate the force required to accelerate the rock.
- **Result:** $F = 50 \cdot 2.0 = 100 \text{ N}$.
- **Instructor Note:** Students see 2,000 N (Rover) vs. 100 N (Rock). This confirms their bias: the "stronger" object has more force and "gives" some to the weaker object.

(Copilot ECR)

Here's a 6-problem sequence all on Mars, built around the "rover vs rock" story. I'll take $g_{\text{Mars}} \approx 3.7 \text{ m/s}^2$.

Each problem is ECR-style (students must explain), and each has an instructor note naming the targeted belief.

([See full response from Copilot](#))

Problem 1 – Rover vs tiny rock, "obvious" outcome

A 900 kg Mars rover is driving at 2.0 m/s on level ground. It hits a loose 10 kg rock lying on the surface. After the brief collision:

- The rock shoots forward and ends up moving at 6.0 m/s.
- The rover slows only slightly, to 1.8 m/s.

(a) Which object's speed changes more during the collision, the rover or the rock?

(b) Explain, in words, why you think that object's speed changes more. You may talk about "how much force" each object has if that's how you think about it.

Problem 2 – Same rover vs embedded boulder, "obvious" outcome

Same rover, same speed: 900 kg rover at 2.0 m/s on level ground.

Problem 3 – Forcing numbers onto the belief (small rock case)

Return to the **tiny rock** collision from Problem 1:

Problem 4 – Center-of-mass check on the same collision

Same tiny-rock collision as Problems 1 and 3:

Problem 5 – Reframing the same collision as mutual interaction

Return again to the tiny-rock collision, but now focus explicitly on the interaction.

Problem 6 – Same scenario, now using equal forces to explain different motion

Considerations

Getting started: What you can do *now*

- Incorporate syllabus statement in your course syllabus
- Have a conversation with students
- Share a poll pre- class to know your students and their use of AI
- Continue AI conversation with students
- Keep current as you can about AI



CTLM Can Support You

- Consultations
- Custom Workshops
- Exploring AI in Teaching - relevant topics
- AI Conversation Sessions
- Monthly AI Newsletter
- Other campus partners: [DoIT](#), [Writing Across the Curriculum](#), and [Libraries](#)

Questions

Thank You

Glad to spend time with you today.

[We welcome your feedback](#)

