TAC Guidance on AI and Physics Instruction

Rationale, Evidence, and Implementation

Purpose and scope

Artificial-intelligence tools are rapidly changing how students interact with information, solve problems, write, and code. These changes have significant implications for physics instruction. The Teaching Activities Committee (TAC) offers this guidance to help instructors, teaching assistants, and students preserve genuine learning while making responsible use of tools that will continue to evolve.

Our central concern is learning. Students develop their understanding of physics by working through problems: choosing principles, testing approaches, making and correcting mistakes, interpreting results, and building physical intuition. AI can support this process. It can also remove precisely the intellectual work through which mastery is formed.[1]

The guiding principle of this document is:

AI should support learning, not substitute for a student's own reasoning.

Our recommendations for students and instructors are found below with a focus on undergraduate and introductory graduate physics courses. A summary is also available on the physics department website.

1. What the evidence suggests—and what it does not yet establish

Research on generative AI in education is still quite new, and conclusions depend strongly on the tool, its safeguards, the students, the subject, and the methodology.[2] The current evidence does not justify either the claim that AI is intrinsically harmful to learning or the claim that ordinary access to a general-purpose chatbot reliably improves learning. It does support several narrower conclusions relevant to physics instruction.

---

[1] Andrew Heiss (2026), "AI for Me, Not (Yet) for Thee? Desirable Difficulties and Deliberate Friction with LLMs." A pedagogical position paper on developing and maintaining expertise, rather than an experimental study of learning outcomes.

[2] Michael Gerlich (2025), "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," Societies 15(1), 6 (correlational study); Hao-Ping Lee et al. (2025), "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers," CHI '25, Article 1121 (self-report survey). These studies raise related questions about cognitive effort but do not establish causal effects on physics learning.

Performing well with AI assistance is not the same as durable long-term learning. This distinction is well established in the learning-science literature.[3] Use of AI can make practice easier or improve immediate performance without producing the relatively lasting changes in understanding that appear when a student must work independently.

A randomized study of nearly 1,000 high-school mathematics students illustrates the danger. Access to a general GPT-4 interface improved performance during assisted practice, but students subsequently performed worse than the no-AI control group on an unassisted assessment. A purpose-built tutoring interface with teacher-designed safeguards improved assisted practice while largely avoiding that later penalty. The result is not directly a study of university physics, but it gives causal evidence for the mechanism faculty are observing: unrestricted answer generation can improve homework while weakening the learning that the assignment was intended to produce.[4]

Conversely, a randomized study in introductory college physics found that a carefully designed AI tutor produced greater learning gains in less time than an in-class active-learning lesson.[5] This is important evidence that AI can support real learning. It should not, however, be generalized to every chatbot interaction. The tutor was deliberately structured around pedagogical practices such as active engagement, graduated assistance, and student explanation. Taken together, the two studies point toward a design principle rather than a simple verdict: how the AI is used matters at least as much as whether it is used.[6]

The evidence base remains limited. Many studies are short-term, involve a single course or subject, and cannot yet tell us how habitual AI use affects the development of independent problem-solving ability over several years. Departmental guidance should therefore be clear but revisable, and course policies should remain tied to explicit learning goals.

2. Recommendations for students

2.1 Your understanding is built by working through problems yourself.

The struggle with a problem is not wasted time. Students who rely on AI to solve homework problems are unlikely to develop the understanding needed for success in physics. Reading a polished solution can create the illusion of understanding, but genuine mastery comes from the process of struggling with a problem: deciding what principle applies, often trying different

[3] N. C. Soderstrom and R. A. Bjork (2015), “Learning Versus Performance: An Integrative Review,” Perspectives on Psychological Science 10(2), 176–199.

[4] H. Bastani et al. (2025), “Generative AI without guardrails can harm learning: Evidence from high school mathematics,” Proceedings of the National Academy of Sciences 122(26), e2422633122. Published version of the 2024 working paper. For an accessible account of that working paper, see Knowledge at Wharton (2024), “Without Guardrails, Generative AI Can Harm Education.”

[5] G. Kestin et al. (2025), “AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting,” Scientific Reports 15, 17458.

[6] Related evidence: M. Lehmann, P. B. Cornelius, and F. J. Sting (2025), “AI Meets the Classroom: When Do Large Language Models Harm Learning?” arXiv:2409.09047v2 (preprint). The experiments found no overall effect on learning; analyses linking outcomes to how students used the tools were exploratory.

approaches, making errors, correcting them, and learning to recognize what a physically reasonable answer looks like.

AI can be an extremely powerful tool for organizing information, explaining ideas, and providing feedback. But using it well requires maturity and judgment.

The best way to develop that judgment is first to work problems yourself, without immediately turning to AI for a solution. Before seeking assistance, clarify what the problem is asking, identify knowns and unknowns, recall relevant principles, attempt a setup, and determine where the attempt fails. A student who can identify a precise obstacle is also in a much better position to obtain useful help from a TA, instructor, peer, or AI tool.

In physics, the process matters. Thinking hard about a problem is what develops intuition.

2.2 If permitted, use AI as a tutor rather than a solver.

If your course permits AI use, consult it only after making a serious attempt on your own. Ask for a conceptual explanation, questions that help you identify where your reasoning breaks down, a limited hint, feedback on your own attempted setup, or a simpler analogous problem. You can also ask it to distinguish a conceptual error from an algebraic one. Do not ask AI to solve an assigned problem or produce a submission for you. Appendix B gives a sample prompt for requesting this kind of tutoring.

Even tutor-style prompting is not a guarantee. A general-purpose model may reveal too much, make a confident error, or gradually drift from the original instruction. Students must remain active participants and should turn to human help when the issue is course-specific notation, an ambiguous problem, or an apparent conflict with the instructor's method.

2.3 Put the tool aside and check your understanding.

After receiving permitted tutoring help, put the tool aside and work through the problem yourself. Explain why the governing principle applies, reconstruct the reasoning, check your calculations using units, signs, limiting cases, estimates, or test cases, and interpret the result physically. If possible, try a short related problem. If you cannot explain every step or solve a closely related problem without assistance, you have not yet finished learning from the interaction.

AI tutoring can include mistakes. Check conceptual explanations and hints against course materials and your own reasoning; if an apparent conflict remains, ask the instructor or TA.

2.4 Follow course rules and retain authorship.

The work you submit must reflect your own reasoning. If your course permits AI tutoring, use it to help you work through problems yourself, not to generate solutions for submission. Follow your instructor's rules on permitted use and disclosure, and be prepared to explain every step without AI assistance. You remain responsible for the accuracy of every claim, calculation, citation, and line of code you submit. Do not upload private information, another student's

work, assessment material that your instructor has prohibited sharing, or copyrighted course material without permission.

3. Recommendations for instructors

3.1 State an explicit AI-use policy and explain why it serves the course.

Students should not have to infer whether AI use is permitted. Every course should state which uses are allowed, restricted, or prohibited; whether material use must be disclosed or cited; and how violations will be handled. Concrete examples are more useful than a general statement that students may or may not use AI. A policy might distinguish between conceptual explanation, brainstorming, hints, checking a student's own work, generating code, editing prose, and producing a complete solution.[7]

The explanation matters. Students may understand that unrestricted AI use can impair learning but still feel pressure to use it because they believe their peers are doing so. A common, clearly explained policy reduces that problem. It also makes the instructor's pedagogical purpose clear: the restriction is not just a defense of older methods, but an effort to preserve the learning the assignment was designed to create.

Instructors should discuss the policy with their TAs before the course begins. TAs need a shared understanding of what they may do with AI in preparing sections or feedback, what students may do, how to respond when students ask for help, and when a possible violation should be referred to the instructor. Instructors should likewise disclose any material use of AI in course materials or feedback.

No single policy will fit every physics course. A prohibition may be appropriate when students are first learning a foundational skill. Conditional use may be appropriate when AI can support a learning goal without simply performing the target skill. More extensive use may be appropriate in an advanced computational course. Whatever the policy, students should not be required to use an external AI service without considering access, privacy, and an appropriate alternative.

3.2 Keep homework central to learning, but pair it with independent assessment.

Homework remains central to physics education. Longer, more challenging problems allow students to integrate ideas, develop technique, and build stamina in ways that short quizzes and exams cannot. The availability of AI does not make that practice less valuable. It does

[7] Inside Higher Ed (2025), "Survey: College Students' Views on AI," Student Voice survey. This reports students' perceptions and preferences, including support for clearer policies; it does not measure learning outcomes.

make the submitted homework score a much less reliable measure of what an individual student understands or can actually reproduce independently.[8]

Faculty report anecdotally that homework grades have risen without comparable improvement in exam performance. Students have always been able to consult friends, solution manuals, or online sources, but general-purpose AI makes polished, individualized solutions immediate and nearly effortless. The appropriate conclusion is not that homework should disappear. It is that homework should be treated primarily as guided practice and formative feedback, with its weight in the final grade calibrated accordingly.

Instructors can reinforce the learning function of homework by requiring students to show their assumptions, strategy, intermediate reasoning, dimensional or limiting-case checks, and physical interpretation. These practices do not make an assignment AI-proof, but they direct attention toward the habits the course values. Feedback should reward those habits, not just the final answer.

3.3 Use multiple, independent demonstrations of understanding.

Because out-of-class work alone cannot reliably establish individual mastery, it should be paired with regular opportunities for students to reason without electronic assistance. Useful formats include:

- short, low-stakes quizzes;
- board work or brief oral explanations;
- guided problem solving during discussion sections;
- structured group work with individual accountability;
- hybrid assignments with an out-of-class component and a short in-person follow-up; and
- occasional requests to explain, modify, or check a submitted solution.

Some quiz questions can revisit a principle or technique used in the homework in a shorter or closely related form. This gives students a concrete reason to work through the homework themselves. It does not mean all homework needs to resemble exam questions: extended and computationally demanding homework serves learning goals that a short in-person assessment simply cannot.

Frequent low-stakes checks are preferable to moving the entire course grade onto one or two high-stakes examinations. They provide students and instructors with earlier evidence of misunderstanding, reduce the penalty of any single poor performance, and encourage continued engagement throughout the quarter.

3.4 Conduct examinations in controlled settings.

Written examinations should be held in class or in another controlled in-person setting without access to AI or other unauthorized electronic resources. Expectations for phones, watches,

[8] A. L. Glass and M. Kang (2022), “Fewer students are benefiting from doing their homework: an eleven-year study,” Educational Psychology 42(2), 185–199. This study of homework and later exam performance predates widespread generative AI use; it is relevant to answer copying, rather than a direct test of AI.

bathroom breaks, and other devices should be made explicit and applied consistently. For small courses, oral examinations or short oral follow-ups can be especially effective in assessing conceptual understanding of course material.

Assessment design should focus on the reasoning students are expected to possess, not on detecting whether a piece of work was produced by AI. Automated AI detectors are not a reliable basis for academic-integrity judgments. A better strategy is to create repeated occasions where students must demonstrate understanding independently and to follow established University procedures when there is concrete evidence of a violation.

3.5 Give task-specific guidance for writing, laboratory work, and code.

Lab reports, papers, presentations, and code raise different questions from problem sets. The instructor should specify the relevant learning goal and then state which stages may involve AI: locating background material, planning, calculation, coding, debugging, drafting, editing, or presentation design.

“Writing is thinking.” Drafting an explanation, choosing what evidence matters, and organizing an argument can clarify scientific understanding. If AI performs those stages too early, a student may submit fluent prose without understanding the underlying argument.[9] Similarly, code that runs is not evidence that a student understands the algorithm, assumptions, or validation. Students should be able to explain their prose and code, verify AI-generated claims and citations against original sources, and test calculations or programs against known cases.

4. Recommendations for teaching assistants

TAs are central to making these policies work in practice. Before teaching begins, TAs should review the course policy with the instructor and agree on how it applies in discussion sections, office hours, grading, and communications with students.

When helping students, TAs should ask to see the student’s attempt and diagnose the earliest conceptual or procedural gap before explaining how to solve the problem. Framing questions, providing limited hints, and giving analogous examples often help the student more than a polished derivation. This is also the behavior students should be encouraged to request from an AI tutor.

TAs should use AI in preparing instructional material or feedback only as authorized by the instructor. They remain responsible for checking all generated content, protecting student information and restricted course materials, and ensuring that students receive meaningful human feedback. Recurrent patterns—declining attendance, widespread inability to explain

[9] N. Kosmyna et al. (2025), “Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task,” arXiv:2506.08872v1 (preprint). This exploratory essay-writing study does not establish long-term effects in physics courses.

submitted work, common AI-generated errors, or confusion about the policy—should be reported to the instructor for an appropriate response.

5. University resources

The Chicago Center for Teaching and Learning provides Guidance for Syllabus Statements on the Use of AI Tools, Designing AI Aware Assignments, Programming and Resources on AI, and Teaching in the AI Landscape. These resources present a range of policy choices rather than assuming that every course should adopt the same rule. Instructors should consult them when drafting course policies and revisit them as University guidance evolves.

Appendix A. Adaptable course-policy language

The following is a starting template, not a department policy requirement:

> The work you submit must reflect your own reasoning. Generative AI tools may be used only as specified by the instructor for the assignment or task. Where AI tutoring is permitted, first make a serious attempt on your own; then use it for conceptual explanations, limited hints, feedback on your attempted reasoning, or a simpler analogous problem. Do not ask AI to solve an assigned problem or generate a solution for submission. After receiving help, put the tool aside and work through the problem yourself. Be prepared to explain every step without AI assistance. For lab reports, papers, presentations, and code, follow the assignment's rules about which stages may involve AI. Disclose material AI use as directed, and take responsibility for the accuracy of all submitted work. Checking or reproducing AI output does not authorize a use prohibited by the course. AI tools are not permitted during quizzes or examinations unless the instructor explicitly states otherwise.

An instructor who prohibits all AI use should say so directly, identify the covered work, explain the learning goals behind the prohibition, and clarify what ordinary tools or human assistance remain available. An instructor who permits broader use should state which uses are permitted, what disclosure is required, and how students will demonstrate independent mastery.

Appendix B. Compact learning-focused tutor prompt

> Act as a learning-focused tutor, not an answer generator. Do not provide the final answer, a complete proof or derivation, finished code, or a hint that effectively reveals the answer.
>
> First ask for the exact problem and instructor directions, my course level, what I think the problem is asking, what I have tried, and the precise point where I became stuck. Diagnose the earliest gap before helping. Begin with the smallest useful intervention: a question, a relevant concept or representation, a subgoal, or a limited hint. Ask only one or two questions at a time. If I make repeated unsuccessful attempts, increase the support gradually; as I improve, return responsibility to me.
>
> Identify the first incorrect or unsupported step and help me repair it. Do not accept a correct answer supported by unexplained reasoning. After I solve the problem, ask me to explain the strategy in my own words, verify the result using appropriate checks, and answer one short transfer question based on the same idea. If the problem is ambiguous or lacks information, say so rather than inventing details. Follow my instructor's notation, permitted methods, and course policy.

Most LLMs have a context window. After a certain amount of information, they can forget earlier parts of a conversation. To reduce this risk, we recommend starting each new problem in a new chat, beginning with the prompt above or something similar.

Acknowledgement: *AI tools were used to help prepare this document. The text has been substantially edited and revised by members of the Teaching Activities Committee.*