


Assessment in an AI-Enabled World: A Faculty Guide

Produced by Northeastern's Task Force on AI and Assessment



Contents

Introduction. →

- 1** | **Begin by retooling one assignment in one course with a careful eye on AI capability.** Ensure that the assignment's assessment strategy validates student learning, deterring inappropriate use of AI, and that the assessment provides credible evidence of what students can do in relation to the learning goals.
- 2** | **Distinguish between practice and assessment.** Make practice activities low stakes, graded on effort or completion rather than correctness. For higher-stakes assessments, use formats where you can confirm that the work was done by students (e.g. work proctored or performed in class).
- 3** | **Factor both process and product into the final grade.** Move away from one-time submissions and toward iterative and incremental assignment activities in which students document stages of their work and development.
- 4** | **Protect the early stages of work.** Structure assignments so that students begin by producing work that demonstrates independent thought and foundational skills in class before they are allowed to integrate AI into their work process.
- 5** | **Make your AI expectations explicit, specific, and explained.** Ambiguity about AI expectations is corrosive. Clearly articulate the boundaries for AI use within the syllabus and at the assignment level, both in writing and verbally in class.
- 6** | **When integrating AI, attend to its framing and positioning.** There is a meaningful difference between using AI to generate a draft that students then edit and using AI to challenge, complicate, or stress-test their thinking.
- 7** | **If you restrict AI, teach students why the restriction matters.** Restrictions are more effective -- and vastly more respected by students -- when they are grounded in a clear explanation of what cognitive work the assignment develops.
- 8** | **Build students' capacity to evaluate AI output in your discipline.** The capacity to assess the accuracy, relevance, and appropriateness of information is a core academic skill. Students need to be able to judge the quality of AI output.

Conclusion/Helpful Links. →

How was this guidance document created? →

Introduction

Something has shifted in higher education, and you've felt it. Student work looks different. Assignments you've relied on for years may no longer measure what you intended them to measure. Colleagues are adapting in different ways, and the guidance you're hearing is sometimes contradictory. If you've been uncertain about where to start, that's not a failure of attention – it reflects the genuine difficulty of a moment that arrived faster than anyone anticipated.

Before turning to practical advice, we want to acknowledge something directly: navigating this moment well requires more than individual faculty effort. The university must continue to serve faculty in robust and concrete ways – through training, tools, infrastructure, policy, and guidance. This document is one part of that commitment, and parallel efforts on infrastructure and policy are underway alongside it.

This guide was written by faculty, for faculty, drawing on discovery work across all ten colleges at Northeastern, six subcommittees, and direct input from students. It is not policy, and it does not ask you to overturn your curriculum. What it offers is practical, evidence-grounded advice to help you ensure your assessments continue to measure what you intend them to measure.

The central finding from that work is this: AI has not primarily created a cheating crisis. AI instead has precipitated a reckoning around the validity of evidence, what we examine to determine whether learning has occurred, and exposed a long-standing weakness in how we assess learning. When AI can produce competent work on a traditional take-home assignment, the assessment instrument – not the student – is the first thing that deserves attention. You know your discipline and what is important, and this reframe ought to drive all subsequent practical decisions about assigning and evaluating student work.

One principle in particular is worth naming here at the outset, because it runs through everything that follows: assessments send a powerful signal to students about what matters in your course. If a grade depends primarily on a final product that AI can generate, students receive the message – whether you intend it or not – that the product is what counts, not the learning behind it. If the grade depends on demonstrated reasoning, iterative improvement, or the ability to explain and defend their work, students receive a very different message. Keeping that signal in mind is one of the most powerful things you can do.

The recommendations that follow are organized to make the work of rethinking your approach to assessment as concrete as possible. They are modest, targeted changes that can serve as a starting point. Each recommendation is grounded in what faculty across this university have already tried, tested, and found useful.

1. Begin by Retooling One Assignment with a Careful Eye on AI Capability

From the outset, keep in mind that you do not need to redesign your entire course at once. Pick one assignment to rethink, for example the assessment you're least confident still measures what it claims to verify.

The most important reframe in this moment is shifting how you name the problem. Faculty across every college have arrived, independently, at the same conclusion: AI is not primarily a cheating crisis. It has exposed the limits of how we have traditionally designed assessments, as evidenced by the ease with which minimally prompted AI can complete many assignments. It's also important to know that detection tools are fallible and not recommended for use by Northeastern.

You know your discipline and what is important, so starting with the following fundamental questions will guide nearly every practical decision that follows:

- I | What cognitive work is your assignment intended to develop and demonstrate? What evidence would demonstrate that a student has actually done that work?
- II | Is the assignment focused on practice, building foundational knowledge that students need to develop independently? Is it intended to help students to apply concepts they've already mastered?
- III | What complex processes will students need to demonstrate, such as analyzing or constructing an argument? Are there additional proficiencies students need to demonstrate, such as creating a specific product or collaborating with others to solve a problem?


Consider what types of student work would provide credible evidence that students have attained the goals for learning (and that cannot easily be offloaded to AI). Faculty report that this kind of clarification often reveals that an assignment has drifted from its original goals for learning. That recognition alone can be a productive starting point for redesign, and therefore it is the place to start (even before considering the assignment's AI policy).



2. Distinguish Between Practice and Assessment

A shift that faculty in predominantly technical and quantitative disciplines have found useful is to separate the practice function of assignments from the assessment function. This is also true for short writing assignments that help students document their thinking and prepare for other coursework. Homework, problem sets, and coding exercises have traditionally served both purposes simultaneously. AI has complicated that dual role, since students can now use AI to complete many take-home tasks without engaging in the cognitive work the tasks were designed to require.

One approach is to make practice activities low-stakes, graded on effort or completion rather than correctness, while reserving higher-stakes assessment for formats where you can be more confident the work reflects the student's own understanding, such as in-class quizzes, exams, or live demonstrations. This doesn't mean practice stops mattering; it means the incentive structure shifts so that practice becomes genuine preparation rather than a grading event that AI can shortcut.



3. Factor Both Process and Product in the Assessment Grade

Ask yourself the following question: If a student used AI to complete this assignment, would they still learn what I need them to learn? If the answer is no, that's a signal the assignment may need redesign. If the answer is yes that may suggest the assignment is already capturing something AI can't easily replicate, or that AI use in that context isn't undermining the learning goal. This is not about making all assignments AI-proof, but to have them be more tightly aligned with the learning you value.

In an AI-ubiquitous world, increasingly the most reliable evidence of student learning in an assessment often lies in the process: the decisions, iterations, and thinking that produced the outcome. In this context, students recommend moving away from one-time submissions and toward iterative and incremental activities. Faculty across colleges are independently moving toward assessing process artifacts, such as drafts, reflections, documented decision-making, version histories, and milestone check-ins.

This is an opportunity to lean into the **experiential learning** potential of the assignment, involving students in real or realistic challenges that provide students with some agency and have value beyond the course. Students report these assignments being extremely valuable and they want more. For help in identifying discipline-relevant assignment ideas, faculty consistently describe peer conversation as the most valuable resource: finding out what a colleague in a related discipline has tried, hearing what worked and what didn't, and discovering that others share the same uncertainties. **AI can also be of assistance in generating ideas**, as you can enter your assignment goals and information about your course into the LLM and prompt it to yield a dozen or so ideas for experiential assignments. Not all AI-generated ideas will be winners, but they can help you brainstorm possibilities.

Process-oriented assessment looks different across disciplines. In writing-intensive courses, it might mean requiring drafts, including student reflections on changes made, to factor the student's decision-making process into the final grade. In technical courses, it might mean asking students to produce reasoning maps or explain their solution path rather than just their answer. In project-based courses, it might mean structured progress checkpoints where students describe what they worked on, what decisions they made, and what changed since the last checkpoint. While it is true that any of the process documents could be generated with AI, including reflections, this approach both deters inappropriate use of AI and sends a signal that process and attention to detail matter in the development of any product.

This approach can be more labor-intensive than grading final products, for both you and your students, so be realistic about what you can sustain. You may want to cut back on other graded assignments to make room for process-oriented work that involves deeper learning. When it comes time to assess the work, keep in mind that you don't need to review every page of exhaustive documentation for every assignment. Selective spot-checking of process artifacts is sufficient. Even better, have students review the assignment rubric and curate their own evidence of learning to provide you with a collection of their work demonstrating their process. This should include a reflective memo in which they tell you what they think matters most and why. This shifts the norm toward transparency and makes it harder for AI-generated work to go undetected.

4. Protect the Early, Generative Stages of Student Work

Across both writing-intensive and technical disciplines, faculty have converged on a key insight: the stages of work most vulnerable to AI bypass are often the most important to protect. The early, generative phases are where much of the learning happens (e.g., exploring a topic, wrestling with a blank page, formulating a thesis, debugging a first attempt). AI is very good at producing adequate first drafts and plausible initial solutions, which means students who start with AI output may never do the foundational thinking those early stages are designed to develop.

You might consider structuring assignments so that students must produce independent work in class before any AI interaction is permitted. This could be a handwritten free-write, a rough thesis statement, an initial attempt at a problem set, or an annotated reading. The goal is to ensure students have something of their own to think with before they encounter AI-generated material. Revision and refinement may be more appropriate stages for AI assistance, once students have generated genuine thinking they can evaluate AI output against.

5. Make Your AI Expectations Explicit, Specific, and Explained

One of the most consistent findings across colleges is that ambiguity about AI expectations is corrosive. Northeastern students report wanting clear boundaries. Communicate your expectations early and often, both in writing and verbally in class.

Consider going beyond a single course-level policy statement. For each major assignment, communicate what AI use is permitted, what is not, and what disclosure is expected. Northeastern's Standards and Recommendations for the use of Generative AI in Teaching and Learning document provides **example policy statements**. Crucially, **explain why** – what knowledge and skills the assignment is designed to develop and how AI use would support or undermine that goal. Students frequently report being more likely to follow restrictions they understand than rules they perceive as arbitrary. You might also consider involving students in discussing the reasoning behind your expectations, which can deepen their understanding of the learning at stake.



6. When Integrating AI, Attend to Its Framing and Positioning

Consider how you position generative AI in the sequence of your assignments. There is a meaningful difference between using AI to generate a draft that students then edit and using AI to challenge, complicate, or stress-test thinking that students have already developed. When AI produces the draft, the core intellectual work – generating ideas, structuring an argument, choosing an approach – has already been done. Students editing an AI draft they didn't conceive often default to surface-level changes rather than substantive engagement. The Digital Education Council's [Next Era of Assessment](#) report provides suggested sequences for AI-integrated assignments.

More productive integrations tend to involve students doing independent work first, then using AI to push that work further: generating counterarguments to a thesis they've already drafted, producing alternative solutions to a problem they've already attempted, or identifying weaknesses in an analysis they've already completed. The key design principle is that AI enters the workflow after the student has something to push back with – not as a substitute for that prior thinking, but as a tool for sharpening it.

7. If You Restrict AI, Teach Students Why the Restriction Matters

Choosing to prohibit or limit AI use in your course is a legitimate pedagogical decision, and you don't need to apologize for it. But restrictions are more effective – and vastly more respected by students – when they are grounded in a clear explanation of what cognitive work the assignment develops and why AI use would bypass exactly that work. The analogy one Northeastern faculty member offered is apt: no one would bring a forklift to lift weights, because the point is to build strength. The case for restriction is strongest when you can specify what strength your assignment is building and why the shortcut prevents it. You might also consider making space to discuss with students what productive struggle feels like and why it matters. Normalizing the difficulty of early-stage intellectual work – rather than treating it as a problem to be solved – can help students understand that the discomfort of not knowing is part of the learning process, not a sign that they need to reach for a tool. Students also express concern about losing out on learning in their efforts to be academically efficient, so joining with them in conversation about discernment around this may be more welcome than you expect.

8. Build Students' Capacity to Evaluate AI Output in Your Discipline

Regardless of where you land on AI use in your course, your students will encounter AI-generated content in their professional lives and will need the ability to judge its quality. Evaluative judgment – the capacity to assess the accuracy, relevance, and appropriateness of information – has always been a core academic skill. AI makes the need for this kind of judgment more visible and more urgent, and students are well-aware of its importance.

Consider creating structured opportunities for students to critically assess AI-generated content in your discipline. In the sciences, AI regularly produces confident but incorrect answers in specialized areas; having students identify where and why those answers fail builds disciplinary knowledge and critical thinking simultaneously. In writing-intensive courses, asking students to compare their own analysis with an AI-generated response – and articulate what the AI misses or gets wrong – develops both subject-matter understanding and metacognitive awareness. In technical fields, students might evaluate whether AI-generated code solves the problem correctly or merely compiles without errors.

This kind of work reframes AI not as a temptation to be resisted or a shortcut to be managed, but as a pedagogical resource for developing the critical thinking your course is designed to build. It also gives students a foundation they will need regardless of how AI policies evolve: the ability to evaluate advice and output from tools of varying reliability, a skill essential to professional practice in every field.



Conclusion

Assessment redesign in response to AI is intellectually demanding work, and it deserves to be treated as such – not as a quick fix, but as a serious engagement with the question of what your students need to learn and how you’ll know they’ve learned it. The good news is that you don’t have to do it all at once, and you already have the expertise that matters most: deep knowledge of your discipline and a commitment to your students’ learning. The guidance above is intended to help you bring that expertise to bear on a genuinely new challenge.

Helpful Links

Northeastern AI Policies

[Standards and Recommendations for the Use of Generative AI in Teaching and Learning](#)

[Policy on the Use of Artificial Intelligence Systems](#)

Teaching Support and Resources

[AI in Teaching and Learning: CATLR sub-site includes a gallery of AI assignments created by NU faculty, tutorials, recordings](#)

[Anthropic Claude courses: Advanced and 101 tutorials + learn new Cowork features](#)

[CATLR website: access workshops, programs, and access to assignment design consultation](#)

[Essential Concepts of AI: Northeastern self-enroll Canvas course](#)

[Assessment in an AI Ubiquitous World: Task Force Briefing Document](#)



How Was this Guidance Document Created?

The guidance above is the product of a thoughtful process that involved input from hundreds of stakeholders. It was developed by Northeastern’s AI in Assessment Task Force. The Task Force, which included faculty representatives from all ten Northeastern Colleges, used the following process in the creation of this document:

- I | Task Force members conducted “discovery scans” within their colleges to identify how faculty are changing their approach to assessment in light of AI, challenges they are encountering, and requests for what would best meet their needs. Results were analyzed to identify high priority needs, with AI assistance.
- II | Six Task Force subcommittees were formed to develop brief reports on priority areas for consideration in AI and Assessment. Subcommittees included Online Asynchronous Courses, Technical Courses, Writing-Intensive Courses, Large Format Courses, Academic Integrity, and AI Grading. Patterns and frequency of themes were identified across the subcommittee reports, with AI assistance.
- III | Northeastern’s Center for Advancing Teaching and Learning Through Research conducted surveys and focus groups to elicit student perspectives on how AI is changing the way they approach their coursework and challenges they are encountering.
- IV | The Task Force synthesized both the subcommittee findings and student perspectives to draft a faculty guidance document. This document was reviewed and revised in response to comments from the Task Force faculty, students, representatives from the offices of the Provost and Chancellor.
- V | The Task Force also generated separate recommendations for consideration related to policies and infrastructure in support of AI and Assessment.

