Protecting Learning:

A Practical Guide to AI-Use Policy for Instructors, Departments, and Schools

Faculty Leads
Augustin Guibaud[1], Manny Patole[2]

Faculty Authors
John Di Bartolo[3], Rosanna Bellini[4], Irene de Lazaro[6], Luke DuBois[7], Semiha Ergan[2,4,5], David Fouhey[8], Peter Li[9], Shizhu Liu[10], Abby Rabinowitz[7,11], Phil Rauscher[12], Jason Röhr[9], Xin Zhang[13]

Student Authors
Samantha Dorminey[14] (graduate), Chaitanya Kukreja[2] (graduate), Yuecheng Shi[4,15] (undergraduate), Isabella Vera[2] (graduate)

Convener
Julia Stoyanovich[4,16]

[1] Mechanical and Aerospace Engineering, NYU Tandon School of Engineering
[2] Center for Urban Science and Progress, NYU Tandon School of Engineering
[3] Applied Physics, NYU Tandon School of Engineering
[4] Computer Science and Engineering, NYU Tandon School of Engineering
[5] Civil and Urban Engineering, NYU Tandon School of Engineering
[6] Biomedical Engineering, NYU Tandon School of Engineering
[7] Technology, Culture and Society, NYU Tandon School of Engineering
[8] Electrical and Computer Engineering, NYU Tandon School of Engineering
[9] General Engineering, NYU Tandon School of Engineering
[10] Technology Management and Innovation, NYU Tandon School of Engineering
[11] Expository Writing Program, NYU Tandon School of Engineering
[12] Chemical and Biomolecular Engineering, NYU Tandon School of Engineering
[13] Finance and Risk Engineering, NYU Tandon School of Engineering
[14] Industrial Engineering, NYU Tandon School of Engineering
[15] Linguistics, NYU College of Arts and Science
[16] Center for Data Science, NYU Courant Institute School of Mathematics, Computing, and Data Science



Fall 2026

NYU TANDON

https://doi.org/10.58153/xp9wz-9gn51

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Preface

In Spring 2026, the NYU Tandon School of Engineering set out to determine what AI means for teaching and for what students learn. The work proceeded from a principle: consequential technology should be governed with the people it affects. Each Tandon department assessed how AI was appearing in its own courses and wrote its own strategy. About 80% of full-time faculty participated in department-level surveys and meetings, documenting what they were observing in their courses and what concerned them.

Three findings from Spring 2026 shaped the work that followed. First, the dominant concern was learning loss, not cheating: students submitting correct work they do not understand. Second, faculty asked for shared guidance they could apply in their own courses: AI-use policy language adaptable for syllabi, ideas for redesigning assignments, and a consistent way to handle suspected misuse. They wanted a common vocabulary and tested language to draw on rather than a mandate applied to every course. Third, many faculty reported that AI is changing both how they teach and what students need to learn: which skills and competencies remain foundational, which are becoming less central, and which new ones are needed.

Over the summer, three faculty working groups addressed what the departments reported: one on AI-use policy, one on measuring how AI is affecting students and their learning, and one on shared teaching resources. This guide, produced by the AI-Use Policy Working Group, is the first to be released.

The guide recommends setting course AI-use policy with students, and two of its authors piloted the approach in Fall 2026. The first pilot ran in two Responsible AI courses: a required undergraduate course of about 190 students and a graduate elective of about 35. Students received a starter policy with the non-negotiable items fixed and explained. Five decisions were open for discussion and a vote: AI use in coding, written work, and course readings; the form of disclosure; and how students would stand behind their work. The second pilot, conducted by one of the working group's leads in a project-based capstone course of about 70 students, used the same design with the open questions adapted to team deliverables and client communications.

Students often chose stricter rules than the instructors expected. The two Responsible AI cohorts resolved the same questions differently, which argues against a single policy for all courses. The instructors adopted the policies the students shaped and will revisit them at mid-semester. These pilots are preliminary. Templates for running the exercise accompany this guide.

This guide was developed through deliberation. Faculty from every department reviewed the evidence, debated the difficult cases, and converged on the positions it presents. Students shaped the pilot policies and are among this guide's authors.

The menus and templates in this guide reflect decisions made at NYU Tandon. The process that produced them transfers better than the decisions themselves. A school that adopts another school's policy without deliberation does not develop its own understanding of the policy. Deliberation also brings faculty and students into direct conversation about learning.

This guide does not represent official NYU policy. It may be adapted for use by other departments, schools, and institutions. Comments and corrections are welcome and may be submitted through the form linked at the end of the guide.

Julia Stoyanovich
Associate Dean of AI Initiatives, NYU Tandon School of Engineering
Brooklyn, New York, October 2026

Preface

Overview: Origin and How to Use This Guide

Why this matters

Core recommendations

Content of the document

How to use this guide

Quick start for instructors: using this guide to update your syllabus

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Overview: Origin and How to Use This Guide

This guide originated as an internal project at NYU Tandon, where the AI-Use Policy Working Group spent the summer of 2026 gathering evidence from departments across the school, reviewing published research, and surveying how other institutions approach generative AI. The resulting guide adapts that work for use by any department, school, or institution building or revising its own AI-use policy. NYU Tandon examples appear throughout as illustrations, not as templates for adoption.

The guide was developed through deliberation. Faculty from every department contributed evidence and reviewed drafts. About 225 students in two pilot courses debated and voted on open policy decisions, and about 70 students in a graduate capstone provided voluntary anonymous input on an adapted version. Student contributors are among the guide's authors.

Reports already exist examining how generative AI is changing teaching and learning; the report of MIT's Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training (Massachusetts Institute of Technology, 2026) is one example. NYU's Office of the Provost and Center for Teaching and Learning publish university-wide AI guidance, teaching resources, and events, which this guide draws on and cites throughout (New York University, n.d.-b). University guidance deliberately leaves the substance of each course's AI-use policy to schools, departments, and instructors: "no central policy can cover the range of possible use cases" (Duff, 2026).

This guide focuses on execution. It consolidates approaches already in use across higher education into a menu of compatible learning objectives, assessment methods, and AI-use policies that a department or an instructor can adopt or adapt. It pairs the menu with practices for shaping the policy with students and a consistent process for handling suspected misuse.

A school is the right unit for this work: small enough to gather evidence from every department and hear directly from faculty and students, large enough that the answers genuinely differ. The framework is not specific to engineering: the evidence was local, but the instrument is designed to transfer. Any school can adapt it or rerun the process for itself.

Why this matters

Generative AI tools can now complete many take-home assignments at a passing level. This changes what a grade certifies and how students engage with coursework outside the classroom. Faculty interviewed across departments agreed that the primary risk is not cheating but learning loss: students submitting correct work they do not understand. In one foundational course, homework scores averaged near 100% while proctored exams on the same material averaged near 50%. AI detectors are unreliable and biased against non-native English writers (Liang et al., 2023), which makes them unsuitable as a basis for enforcement.

A course AI-use policy establishes how AI may be used in a course, ties that use to the learning objectives, and communicates those expectations to students. No single policy fits every course because appropriate AI use depends on what each course is meant to teach. A common vocabulary, tested language, and a fair process can be shared. This guide provides them.

Preface

Overview: Origin and How to Use This Guide

Why this matters

Core recommendations

Content of the document

How to use this guide

Quick start for instructors: using this guide to update your syllabus

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Core recommendations

- **Set a baseline, then tune**. Set a course-wide baseline using the Red (AI Prohibited), Yellow (Limited AI), and Green (AI-Integrated) labels, then adjust it at the assignment level. State the reasoning in the syllabus so that each rule is tied to what the assignment is meant to assess.
- **Expect a profile, not a single color.** Most courses will end up with a profile: Green on coding, Yellow on written analysis, Red on quizzes. The labels make that profile legible to students, but the policy is the decisions behind them: what a disclosure must contain, and how students stand behind their work.
- **Shape it with your students.** Bring a starter policy with the non-negotiables explained, put the genuinely open decisions to your students through discussion and a vote or other structured input, and keep the final call. When we did this in three courses, students engaged seriously, often chose stricter rules than we expected, and consistently preferred accountability to surveillance.
- **Verify mastery AI-free.** Wherever independent mastery matters, add an AI-free check: an oral question, an in-class exam, a proctored quiz. Graded take-home work alone can no longer reliably demonstrate whether a student has learned. Keep disclosure requirements, but never treat them as sufficient, and never treat AI-detector output as proof of a violation.

Preface

Overview: Origin and How to Use This Guide

Why this matters

Core recommendations

Content of the document

How to use this guide

Quick start for instructors: using this guide to update your syllabus

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Content of the document

Section	What it covers
Preface	Why this work was done and how we went about it.
Overview: Origin and How to Use This Guide	The core recommendations, a guide-map by reader, and the instructor Quick Start.
1. Landscape & Scoping Analysis 1.1 Mapping current practice 1.2 What's working, and what isn't 1.3 What published research says 1.4 What this means for the rest of this guide	What a landscape scan of department-level practice found: what is working and what isn't in the classroom, illustrated with findings from one engineering school, and what published research on AI and pedagogy says.
2. Building Your Course Policy 2.1 What every policy should do 2.2 The three levels of AI use 2.3 The customization menu	What every policy must do, the three levels of AI use with real examples, and a clause-by-clause menu for building course and assignment language.
3. Student-Facing Socialization 3.1 Regardless of which policy you adopt 3.2 Shaping the policy with your students 3.3 A student-facing checklist and templates 3.4 Examples: Writing clear AI instructions	How to introduce your policy, shape it with students (two case studies), revisit it during the term, and write clear assignment instructions and disclosure statements.
4. Academic-Integrity Process 4.1 Why AI misuse is not plagiarism 4.2 Relation to existing codes of conduct 4.3 Instructor guidance 4.4 Proposed wording updates 4.5 What peer institutions do 4.6 Open decisions	How AI-related misconduct fits (or doesn't) into a typical Code of Conduct, how instructors can handle suspected cases, proposed wording updates, approaches at peer institutions, and key decisions for schools and departments.
Conclusion	What this guide provides and what remains open, a transparency note on our own AI use, and full references, including links to every cited institutional policy.
References	
Appendix A Appendix B Appendix C Appendix D Appendix E	Policy examples reviewed for this guide. Institutional guidance on shaping policy with students. Components of a worked course profile. AI-use policy starter draft for deliberation. A worked Code of Conduct review.

Preface

Overview: Origin and How to Use This Guide

Why this matters

Core recommendations

Content of the document

How to use this guide

Quick start for instructors: using this guide to update your syllabus

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

How to use this guide

- **If you teach a course:** Section 2 gives you ready-to-adapt policy text built on the Red/Yellow/Green model, Section 3 shows how to introduce the policy to students or shape it with them, and Section 4 tells you what to do if you suspect misuse, and why AI detectors alone are not enough. The Quick Start list sets up a working policy in five steps.
- **If you lead a department or program:** Section 1 shows what a landscape scan of course-level practice finds and how to replicate it; Section 2.3 helps coordinate policy choices across courses; Section 4.6 lists the decisions that need an answer above the course level.
- **If you set policy for a school:** See Section 1 for the evidence base, Section 4.5 for how peer institutions handle suspected misconduct, Section 4.6 for the open decisions, and the Conclusion for considerations after adoption.

Quick start for instructors: using this guide to update your syllabus

If you only have ten minutes, use this checklist; each step points to the section with the full detail, examples, and rationale.

1. **Pick your course baseline.** Choose Red (AI Prohibited), Yellow (Limited AI), or Green (AI-Integrated) as your course-wide default. (Sections 2.2-2.3)
2. **Set a policy level for each assignment type.** Pick a policy level (AI Prohibited / Limited AI / AI-Integrated) and the clauses that apply. (Section 2.3, using either the full customization menu or the faster "à la carte" chart at the end of that section)
3. **Write the assignment-specific language.** Paste the assignment box template into your syllabus or assignment instructions and fill in what's allowed, what's not, whether disclosure is required, and a one-sentence rationale. (Section 3.3; worked examples in 3.4.)
4. **Introduce the policy in Week 1.** Explain the policy and the reasoning behind it during the first week of class, and consider putting the open decisions up for discussion with your students. (Sections 3.1 and 3.2)
5. **Revisit your policy mid-semester.** Set a reminder to bring the policy back up before major assignments and again around midterms; reflect on it with students at the end of the term (Section 3.1).

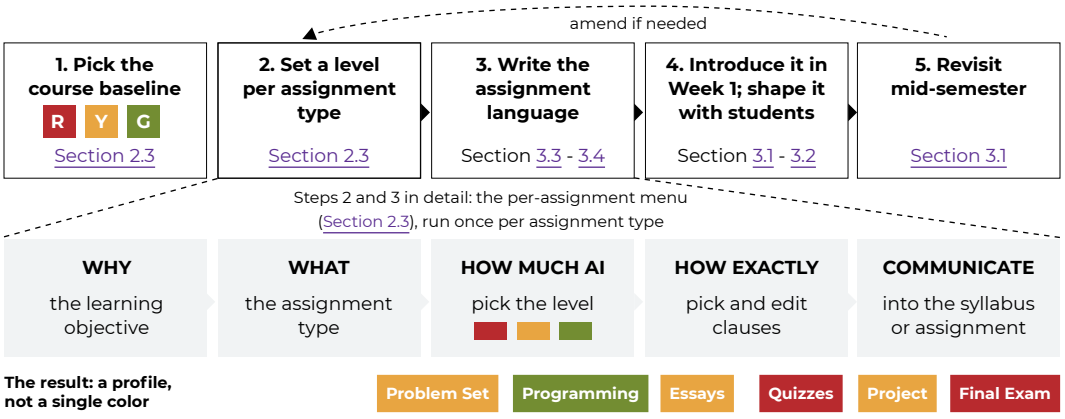


Figure 1. From learning objectives to a working course policy. Top: the five Quick Start steps, with the sections that detail each. Middle: the per-assignment menu of Section 2.3, which expands Steps 2 and 3. Bottom: the output is a profile across assignment types, not a single color. Figure generated with Claude Opus 5.5.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

1.1 Mapping current practice

1.2 What's working, and what isn't

1.3 What published research says

1.4 What this means for the rest of this guide

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

1. Landscape & Scoping Analysis

This section is organized around three questions any institution can ask of itself:

1. What policies and strategies are departments currently using?
2. What approaches are or are not working?
3. What does emerging research on AI and pedagogy say about those approaches?

We illustrate the first two with findings from a structured evidence-gathering exercise at NYU Tandon: a worked example of what such a scan can surface and a starting point for comparison with your own departments. The section closes with what each finding means for the rest of this guide.

1.1 Mapping current practice

Readiness varies widely, and that variation is itself useful information. We mapped current practice across departments, drawing on their AI-strategy submissions. A few departments already had detailed, field-tested policy language: named tiers, a written logic-audit, in which students justify their AI-assisted reasoning, required for medium-trust AI use, or a disclosure model refined over several semesters. Most had informal, instructor-by-instructor norms rather than a written department policy, and a few had not yet formulated any guidance at all. This guide is therefore built as a menu departments and instructors can adopt at their own pace (Section 2).

Assessment has already shifted toward AI-free formats. Nearly every department had shifted weight toward proctored, in-class, oral, or no-device assessment, regardless of its overall stance on AI use. In many courses, this functions as the de facto policy even where nothing is written down. One foundational course provides a clear example: homework averaged near 100% while proctored exams on the same material averaged near 50%. The gap is consistent with AI-supported performance masking deficits in understanding rather than closing them (an "illusion of competence"), and it mirrors findings at other institutions. Both the shift and the underlying gap are well-supported by published research (Section 1.3). There is, however, an important limitation to this approach: securing assessment is not the same as protecting learning. A proctored exam can confirm a student can perform unaided; it does not compensate for the loss of practice and productive struggle when AI substitutes for parts of the learning process.

1.2 What's working, and what isn't

Faculty's primary concern is learning loss, not cheating. When asked what concerned them most, faculty in every department that collected feedback said the same thing, and in some departments' surveys, concern about learning loss outpaced concern about cheating by 30 to 40 percentage points. That gap is why this guide treats learning loss as the central problem an AI-use policy needs to solve, and why policy design (Section 2), student-facing communication (Section 3), and academic-integrity process (Section 4) are treated as three related but distinct levers.

Task-specific AI-use rules allow greater flexibility within courses. Departments that split AI-use rules by task rather than applying one rule to an entire course reported it as useful, under several different names: one department used a red/yellow/green stop-light model; another, three assessment contexts tied to risk level; a third, a four-stage

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

1.1 Mapping current practice

1.2 What's working, and what isn't

1.3 What published research says

1.4 What this means for the rest of this guide

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

adoption-direction framework to manage and measure how instructors integrate AI. Faculty preference data points the same way: in one department's survey, a majority of faculty wanted shared guidance on AI-use policy, but explicitly did not want a single uniform policy imposed on every course. Section 2 is built around this logic: a course-wide baseline, adjusted assignment by assignment.

Faculty already write usable guidance on their own, with or without coordination. In one department, most instructors had already added some form of AI-use language to their syllabi independently, before any coordinated department policy existed. Without coordination, however, language and expectations can diverge across a program. Whether students actually notice or remember scattered syllabus language is a separate, open question to which we return in Section 3.

Disclosure alone is insufficient, and suspected-misuse cases are the clearest common difficulty. Even the most developed disclosure model, which required students to document AI use on every assignment, produced suspected undisclosed use. AI capabilities are now embedded directly into everyday writing and search tools, making AI use more difficult to track. Handling a suspected case after submission was the most consistently reported challenge: in one department's survey, roughly two-thirds of faculty said they had encountered a confirmed or suspected case of unauthorized AI use, and the existing escalation process was widely described as burdensome, with inconsistent outcomes. That finding is why this guide gives process its own section (Section 4) rather than treating policy language as sufficient on its own.

1.3 What published research says

We checked the main practices identified above against peer-reviewed and preprint research available as of summer 2026. Moving high-stakes assessment into the classroom is strongly backed by the literature; task-differentiated ("tiered") policies are promising but supported by limited evidence; and disclosure of AI use is where careful research complicates the intuitive approach. Given the novelty of AI in education, there is not yet a large body of authoritative research, and the education-research literature should be rechecked regularly as it matures. We examine each practice in turn.

Practice 1: Moving assessment into class (proctored, oral, no-device)

The "illusion of competence" pattern flagged in Section 1.1, reflected in high take-home scores alongside a much lower proctored-exam average, is corroborated by controlled research. Fan et al. (2025) found that students supported by ChatGPT on a writing task produced higher-quality work but did not learn more, showing "metacognitive laziness." Rojas-Galeano (2025) found that programming students who reported high understanding during an AI-assisted phase performed significantly worse when asked to extend their code unaided. The pattern holds at scale: Strömberg et al. (2026) tracked 26,811 Chinese secondary students and found that AI adoption raised homework scores by 18% but cut monthly exam scores by 20% and entrance-exam scores by 18-24%, with losses concentrated among students who used AI to shortcut homework rather than engage with it. Using EEG, Kosmyna et al. (2025) found that ChatGPT-assisted essay writers had weaker brain connectivity and memory encoding despite scoring well, and mostly could not quote their own writing.

Securing high-stakes assessment through in-person, oral, or proctored formats is now standard advice at peer institutions, though the research flags real costs: oral verification is resource-intensive, inconsistent across markers, and raises anxiety and fairness

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

1.1 Mapping current practice

1.2 What's working, and what isn't

1.3 What published research says

1.4 What this means for the rest of this guide

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

concerns. Goodhart's Law, in Strathern's (1997) phrasing, holds that "when a measure becomes a target, it ceases to be a good measure." Applied here, assessment can verify whether learning occurred but cannot substitute for the practice, feedback, and productive struggle through which learning develops. Heavy reliance on AI may interfere with that process: Choudhuri et al. (2026) found that reflection and critical thinking decline most sharply among students with the most confidence in the tool, contributing to a cognitive-debt cycle. An AI-free check is thus a necessary complement to a policy, not a complete policy on its own (Section 2).

Practice 2: Task-differentiated ("tiered") policies

The tiered logic that emerged independently across several departments in our own scan (Section 1.2) mirrors the AI Assessment Scale (AIAS) of Perkins et al. (2024), a five-level scale from no AI use to full integration, typically applied per assignment or task rather than as a single course-wide setting. A pilot implementation (Furze et al., 2024) reported fewer generative AI-related misconduct cases and improved engagement after adoption. The evidence remains limited, however: the pilot covers a single institution, not a controlled comparison, and the framework's authors have since published a corrective on common misapplications (Perkins et al., 2025), such as over-specifying AI-use levels without redesigning the underlying task.

Practice 3: Disclosure of AI use

The disclosure difficulties described in Section 1.2, where even a well-developed, several-semesters-old disclosure model still produced suspected undisclosed use, are well documented in the literature. At King's Business School, up to 74% of students across undergraduate and postgraduate modules failed to complete a required AI-use declaration (Gonsalves, 2025). The same study identifies several barriers to compliance, including fear of academic repercussions, ambiguity about what must be declared, inconsistent enforcement, and peer influence, and recommends clearer and more consistent guidance, supportive communication, and greater institutional trust rather than abandoning disclosure requirements.

Schilke and Reimann (2025) studied disclosure beyond academia. Across 13 experiments in a range of professional contexts, they found that people who disclosed AI use were consistently trusted less than those who did not. The effect persisted across framings and under both voluntary and mandatory regimes; at the same time, third-party exposure of previously undisclosed AI use produced a larger trust penalty than self-disclosure. This research highlights a real social cost of disclosure, to which students entering the job market will respond. The fear of being trusted less for being honest about AI use may help explain the compliance gap Gonsalves documents.

Across these three practices, the evidence provides varying levels of support. In-class and proctored assessment is the best supported: the research validates the take-home/proctored-exam gap identified in Section 1.1, and the practice verifies independent capability, not AI skill. Tiered, task-differentiated policies are promising and externally validated as a framework, but the evidence behind them is convergent practice plus one pilot study; calling them proven would require outcome data that does not yet exist. The evidence for disclosure is the weakest: it supports transparency, but no study shows it reliably ensures compliance or demonstrates learning, so its value depends heavily on how it is defined, communicated, and enforced.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

1.1 Mapping current practice

1.2 What's working, and what isn't

1.3 What published research says

1.4 What this means for the rest of this guide

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

1.4 What this means for the rest of this guide

- **For Section 2 (Building Your Course Policy):** use a course-wide baseline with assignment-specific categories, tie each rule to the relevant learning objective, and pair it with an AI-free check wherever independent performance must be demonstrated.
- **For Section 3 (Student-Facing Socialization):** present the policy as part of teaching and learning, not only as a warning about misconduct: explain the reasoning, revisit it during the term, clarify assignment-level expectations, and treat disclosure as a teaching tool rather than an enforcement mechanism.
- **For Section 4 (Academic-Integrity Process):** handle suspected AI misuse with clear definitions, explicit course- and assignment-level rules, an initial conversation with the student, multiple forms of evidence, and caution about relying on AI-detector results alone.
- **For your institution's process:** treat the landscape scan in Sections 1.1 and 1.2 as a starting template, not a one-time exercise. The landscape is moving quickly; repeat the scan, or build a standing channel for departments to report what they are trying, and track how your own policies perform once adopted (see the Conclusion).

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

2. Building Your Course Policy

This section walks the instructor from requirements to finished syllabus language: what every policy should do regardless of stance (2.1), the three levels of AI use that are its building blocks (2.2), and a customization menu that assembles levels and clauses into assignment-specific language (2.3). Consistent with the findings of Section 1, everything here assumes a course-wide baseline with assignment-level adjustment, not a single mandatory rule. Real institutional policies are collected in Appendix A; examples of clear assignment wording are in Section 3.4.

2.1 What every policy should do, regardless of stance

1. State clearly what students may and may not do.
2. Explain how the rule supports the learning objectives of the course or the assignment.
3. State what record of AI use, if any, is required: a disclosure, a citation, saved prompts and outputs, or another process record.
4. Protect personal, confidential, proprietary, assessment, and unpublished material.
5. Use AI-free quizzes, exams, oral checks, or live work when independent mastery matters.
6. Tell students what rule applies when an assignment is silent, and whom to ask when uncertain.
7. Work for all students: when restricting tools, say whether grammar, spelling, and translation aids count, and account for multilingual writers; approved accommodations always remain available.

2.2 The three levels of AI use

Every policy in this guide is built from three levels, and most courses will use more than one.

Level	What it means	What it looks like in practice
Red (AI Prohibited)	Students complete the specified work without generative AI, because independent performance is part of the learning objective.	Paired with in-class or proctored assessment; chosen where a specific skill must be built without assistance (a writing program at Stanford; a systems course at NYU Courant).
Yellow (Limited AI)	Students may use AI only for the listed tasks, tools, or stages; all other use is prohibited.	The boundary is stated per assignment: checking work, explaining concepts, or providing feedback, but not producing the submission (NYU Courant Mathematics; Harvard's mixed template).
Green (AI-Integrated)	Students may use AI for the stated work, disclose material use, verify the output, and remain responsible for the final submission.	AI use is expected and taught, and mastery is still verified AI-free (a data-structures course at NYU with AI as tutor and programming partner, and AI-free exams).

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Course AI-use policies now exist at every level of permissiveness, set variously at the course, program, and school level, and there is no consensus among them; the right level depends on what a course is meant to teach. Appendix A lists the course, program, and teaching-center policies we reviewed, with links. We cite them to illustrate the range of existing approaches. This guide recommends choosing a level per assignment (Section 2.3), not one policy for a whole course, department, or school. In our own two-course pilot, students given these choices did not pick one color: they arrived at a profile across assignment types (Section 3.2).

2.3 The customization menu: from learning objectives to syllabus language

It is unlikely that a course requires a single, uniform AI-use policy across the full range of its activities. Set a course-wide baseline using the levels of Section 2.2, then adjust it assignment by assignment. If you arrived here from the Quick Start, this menu provides Steps 2 and 3 in detail, applied to each assignment type.

Each assignment should state: allowed and prohibited tasks; approved tools and protected-material restrictions; required disclosure, citation, or process record; any AI-free mastery check; whether the assignment overrides the course baseline; and the rationale, one sentence tying the rule to the learning objective.

> *Simple disclosure template: "I used [tool/version] for [purpose]. I checked [facts, citations, calculations, or code] and remain responsible for the final submission."*

We offer here a granular tool for instructors to build assignment-specific clauses rather than start from a single template. The menu works with the three levels of Section 2.2 and follows five steps, and examples are provided below.

Step	Faculty Decision
1. WHY?	Select the learning objective(s) being assessed.
2. WHAT?	Select the assignment type and subcategory.
3. HOW MUCH AI?	Select AI Prohibited, Limited AI, or AI-Integrated.
4. HOW EXACTLY?	Select, edit, or delete applicable clauses from the policy customization modules.
5. COMMUNICATE	Paste the resulting language into the syllabus and assignment instructions, labeled: Red (AI Prohibited), Yellow (Limited AI), or Green (AI-Integrated).

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Step 1: Learning objectives for each assignment

Start from what each assignment is meant to assess. For the selected assignment, choose the primary learning objective(s):

Learning Objective	Faculty Decision	Potential AI Integration	Restricted AI Usage
Independent Mastery	Independently demonstrate the knowledge and skills being assessed.	AI may support learning and preparation.	AI use is prohibited whenever independent performance is the learning outcome.
Knowledge Acquisition & Conceptual Understanding	Build foundational knowledge and accurate mental models.	Explain concepts, provide examples, generate practice.	Does not replace student responsibility for independent understanding.
Analytical, Critical & Quantitative Reasoning	Develop logical reasoning, mathematics, analysis, and judgment.	Check work, explain methods, compare approaches.	Does not replace student responsibility for independent skill demonstration.
Problem Framing & Systems Thinking	Define problems and understand interconnected systems.	Brainstorm problem statements and system relationships.	Does not replace student responsibility to independently justify the problem definition.
Engineering Design & Creativity	Develop original engineering and design solutions.	Generate ideas and iterate designs.	Students retain ownership of design decisions and reasoning.
Programming, Data & Computational Thinking	Develop coding, debugging, modeling, simulation, and data analysis skills.	Explain code and assist analysis.	Students remain responsible for demonstrating independent computational competence.
Research & Information Literacy	Locate, evaluate, and synthesize evidence.	Assist literature discovery and organization.	Does not replace engagement with and verification of original sources.
Communication & Professional Practice	Develop communication, collaboration, ethics, and professional judgment.	Improve clarity, organization, and documentation.	Students remain authors of ideas and recommendations.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Step 2: All course assignments

There is a broad range of assignments that can be offered as part of a course. This approach can be used across all of them. As an illustration, the following table summarizes the assignment taxonomy at NYU Tandon:

Homework	Quizzes	Exams	Group Projects	Individual Projects	Papers
· Problem Sets & Analytical Exercises · Programming & Computational Assignments · Laboratory & Data Assignments · Reading & Written Response Assignments · Design & Creative Exercises	· Objective & Short-Answer Quizzes · Quantitative & Analytical Quizzes · Programming & Technical Quizzes · Practical, Laboratory & Oral Quizzes · Take-Home / Web-Based Quizzes	· Written & Quantitative Exams · Programming & Computing Exams · Case-Based & Design Exams · Practical, Laboratory & Oral Exams · Take-Home & Open-Resource Exams	· Engineering Design & Development Projects · Programming, Computing & Data Projects · Research, Laboratory & Analytical Projects · Design, Media & Creative Projects · Business, Policy & Professional Projects	· Engineering Design & Development Projects · Programming, Computing & Data Projects · Research, Laboratory & Analytical Projects · Design, Media & Creative Projects · Capstone, Thesis & Professional Projects	· Analytical & Reflective Papers · Research Papers & Literature Reviews · Technical & Engineering Reports · Project Reports & Proposals · Professional, Policy & Business Writing

Step 3: Select the AI policy level for each assignment

Select one policy level as the baseline for an assignment. Instructors can select one or more of the policy statements associated with the chosen level. Ensure that the selected statements are consistent with one another.

AI Prohibited	Limited AI	AI-Integrated
Examples: ☐ Generative AI may not be used in any form. ☐ AI may not generate, rewrite, summarize, or edit your responses. ☐ All submitted work must be entirely your own. ☐ Use of AI will be considered unauthorized assistance. ☐ LockDown Browser or other secure technologies may be required when specified.	*Examples:* ☐ AI may assist with brainstorming, organizing ideas, or improving clarity. ☐ AI may be used only for proofreading grammar, spelling, or clarity. ☐ AI may not generate substantive content for your responses. ☐ AI-generated content must be reviewed, verified, and substantially revised before submission. ☐ You remain responsible for every submitted statement. ☐ Provide a disclosure describing any permitted AI use if requested. ☐ You may be asked to explain your responses without AI assistance, or provide your prompt history.	*Examples:* ☐ AI use is expected as part of completing this assignment. ☐ AI may assist with drafting, revising, summarizing, or refining responses. ☐ Students are expected to supervise AI output rather than simply accept it. ☐ Students remain fully responsible for all submitted work regardless of AI assistance. ☐ Document the AI tools, their purpose, and the extent of their contribution. ☐ Submit required prompts, interaction logs, or other process documentation if specified. ☐ You may be required to reflect on how AI influenced your reasoning or decision-making.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Step 4: Assignment instruction customization

Once the policy level and policy statements for an assignment have been selected, specific instructions may be more restrictive or more permissive than the label alone indicates. The options below are designed to cover a range of AI-use instructions identified in this study.

Policy Module	Policy Customization Options
1. Permitted AI Functions	☐ Brainstorm ideas, questions, approaches, or design alternatives. ☐ Explain concepts, methods, terminology, or course content. ☐ Generate examples or practice problems that are not submitted for credit. ☐ Outline or organize the structure of work. ☐ Proofread grammar, spelling, punctuation, or formatting. ☐ Improve clarity, readability, or tone without changing substantive meaning. ☐ Check completed calculations, reasoning, or answers after an independent attempt. ☐ Provide feedback on student-created work. ☐ Debug student-written code or identify possible coding errors. ☐ Assist with data cleaning, analysis, visualization, modeling, or simulation. ☐ Generate drafts or portions of submitted written work. ☐ Revise, rewrite, or summarize substantive content. ☐ Generate code or formulas. ☐ Generate designs, images, audio, video, or other submitted content.
2. Prohibited AI Functions	☐ Generate final answers or responses for submission. ☐ Solve assigned problems or complete assessed reasoning on the student's behalf. ☐ Generate substantive written content submitted for credit. ☐ Generate code, calculations, models, or designs submitted as the student's own work. ☐ Rewrite student work in a manner that replaces the student's authorship. ☐ Summarize assigned readings in place of engaging directly with the original material. ☐ Generate research findings, experimental observations, interview responses, or other evidence. ☐ Generate, fabricate, alter, or conceal data. ☐ Invent citations, quotations, sources, results, or factual claims. ☐ Upload course, assessment, partner, proprietary, confidential, personal, or unpublished research materials without authorization.
3. Disclosure and Attribution Requirements	☐ No disclosure is required for the permitted use. ☐ State that AI was used. ☐ Identify the AI tool and, when relevant, the version used. ☐ Describe the purpose for which AI was used. ☐ Describe the extent of AI's contribution. ☐ Identify the problems, sections, code, figures, or other portions influenced by AI. ☐ Distinguish student-created content from AI-generated or AI-revised content. ☐ Attribute or cite AI-generated material when required by the instructor or applicable citation style.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Policy Module	Policy Customization Options
4. Process Documentation and Records	☐ No AI process records are required. ☐ Submit the prompts used. ☐ Submit relevant AI outputs. ☐ Submit the complete conversation or interaction history. ☐ Include an AI-use statement or appendix. ☐ Retain and provide drafts, notes, or version history when requested. ☐ Retain and provide code repositories, commit logs, or change histories when requested. ☐ Record which AI suggestions were accepted, rejected, or modified. ☐ Preserve relevant interactions for possible instructor review.
5. Verification and Student Responsibility	☐ Verify all factual claims against credible sources. ☐ Verify all quotations, citations, references, and source details. ☐ Verify calculations, formulas, assumptions, and units. ☐ Test AI-generated or AI-modified code for correctness, security, and intended behavior. ☐ Validate data analysis, visualizations, models, simulations, and statistical results. ☐ Check for bias, omissions, unsafe recommendations, or inappropriate conclusions. ☐ Identify and correct AI errors before submission. ☐ Ensure the work complies with course instructions, academic-integrity requirements, professional standards, safety requirements, and applicable laws or policies. ☐ Remain fully responsible for the accuracy, quality, originality, and appropriateness of all submitted work.
6. Independent Mastery and Verification	☐ Attempt the work independently before consulting AI. ☐ Document the independent attempt made before AI was used. ☐ Be able to reproduce the work without AI assistance. ☐ Be able to explain and defend every submitted answer, design choice, calculation, source, and line of code. ☐ Complete an AI-free quiz, exam, oral check, presentation, demonstration, or code review related to the assignment. ☐ Respond to targeted follow-up questions about AI-assisted work. ☐ Revise or extend the work under supervised or AI-free conditions. ☐ Demonstrate individual competence when work was completed collaboratively or with AI assistance.
7. Reflection and Critical AI Literacy	☐ Explain where AI was helpful and why. ☐ Identify where AI was inaccurate, incomplete, biased, or misleading. ☐ Compare the AI-generated approach with the student's independent approach. ☐ Explain why specific AI recommendations were accepted, rejected, or modified. ☐ Evaluate the limitations, assumptions, uncertainty, and risks of the AI output. ☐ Reflect on how AI affected the student's reasoning, creativity, efficiency, or learning. ☐ Assess whether AI use improved or weakened the final work. ☐ Describe how the student maintained authorship, professional judgment, and accountability.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Policy Module	Policy Customization Options
8. Privacy, Security, Access, and Tool Restrictions	☐ Use only instructor-approved AI tools. ☐ Use only institutionally provided AI services when specified. ☐ Do not upload copyrighted course materials, assignment prompts, answer keys, examinations, or other protected assessment content without authorization. ☐ Do not upload personal, confidential, partner, proprietary, export-controlled, regulated, or unpublished research information without authorization. ☐ Follow institutional data-classification, privacy, cybersecurity, research, and records requirements. ☐ Use an institutionally provided or free tool when AI use is required, or use the equivalent alternative provided by the instructor. ☐ Approved disability accommodations remain available. ☐ Do not create an account, accept external terms, or provide personal data when an equivalent course-approved alternative is available. ☐ Ask the instructor before using an unlisted tool or feature.

Step 5: Communicate

The following worked examples can be used as reference points when developing or revising the syllabus and assignment guidelines.

> ***Take-Home Short-Answer Quizzes (15% of grade)*** - *"AI Prohibited" "Complete all quiz responses independently, without generative AI assistance. AI may not be used to interpret questions, identify answers, generate or revise responses, or summarize course materials. The purpose of this assignment is to assess your independent understanding and application of course concepts. You must be able to explain and support all submitted responses."*

> ***Homework Problem Sets (15% of grade)*** - *"Limited AI" "AI may be used to explain concepts, provide practice examples, or check your solutions after you have independently attempted the assigned problems. AI may not be used to solve assigned problems or generate answers for submission. Identify which problems you used AI assistance on. You remain responsible for verifying all calculations and must be able to explain and reproduce your submitted solutions without AI assistance."*

> ***Monte Carlo Group Project (20% of grade)*** - *"AI-Integrated" "You may use AI to support brainstorming, organizing ideas, programming assistance, debugging, data analysis, technical writing, and improving clarity. AI may contribute to submitted work. You must review, verify, and substantially revise all AI-generated content before submission and remain fully responsible for accuracy, completeness, and appropriateness. Disclose all AI tools used and how they contributed; include prompts or process records if requested. Teams may be asked to explain or reproduce any portion of the project without AI assistance."*

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

2.1 What every policy should do, regardless of stance

2.2 The three levels of AI use

2.3 The customization menu: from learning objectives to syllabus language

Step 1: Learning objectives for each assignment

Step 2: All course assignments

Step 3: Select the AI policy level for each assignment

Step 4: Assignment instruction

Step 5: Communicate customization

An "à la carte" syllabus summary

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

An "à la carte" syllabus summary

The menu's output is a per-assignment summary the instructor pastes into the syllabus, so students see the whole profile at a glance. A worked excerpt:

Assignment Type	AI Prohibited	Limited AI	AI-Integrated
Homework: Math Problem Sets			AI may be used to help you learn how to solve the problem. Disclose material use.
Take-Home Short-Answer Quizzes (platform-based)	Complete quizzes independently; AI should not be used to interpret questions, generate responses, revise answers, or locate information.		
Group Projects: Capstone		Use AI only for designated project tasks; disclose assistance, verify outputs, and maintain team responsibility*.	

* Note that AI-detection tools are not generally reliable, with high false-positive rates. The results of such tools should not be considered definitive evidence of AI use or misuse in academic code-of-conduct investigations. See Section 4 for further discussion.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

3.1 Regardless of which AI policy you adopt

3.2 Shaping the policy with your students

Case study 1

Case study 2

3.3 A student-facing checklist and copy-ready templates

Suggested assignment box

Suggested student AI disclosure statement

3.4 Examples: writing clear AI instructions

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

3. Student-Facing Socialization

A course AI-use policy only works if students know it, understand it, and understand its rationale. Many students already use AI tools daily and are uncertain about what is allowed in your course; some do not realize that inappropriate use can constitute academic misconduct. As with any classroom policy, communicate it clearly, explain why it is in place, and treat it as part of teaching the course rather than as fine print.

A policy is most effective when it does more than state what is permitted and prohibited: tie it explicitly to the learning objectives of the course and its assignments, and explain how each expectation supports the knowledge and skills being developed. Framing the policy in relation to learning objectives makes clear how its requirements are intended to support learning.

This section covers how to introduce the policy you chose in Section 2, how to shape parts of it with your students, how to revisit it during the semester, and how to write assignment instructions specific enough that a student can follow them.

3.1 Regardless of which AI policy you adopt

- ➜ Include your AI-use policy in your course syllabus and post it on the course learning-management system (LMS); discuss it during the first week of class.
- ➜ Explain the reasoning behind the policy; connect it to the learning objectives of the course or a given assignment.
- ➜ Clearly state expectations for attribution, citation, or disclosure of AI use, and state whether rules differ across in-class work, homework, quizzes, and the final exam.
- ➜ Disclose your own use of AI in the course (materials preparation, feedback), and say what you will not do, such as upload student work to AI tools without consent, treat AI-detector output as evidence of cheating on its own, or present AI-generated feedback as your own.
- ➜ Give students opportunities to ask clarifying questions; keep a running FAQ from those questions and add it to the syllabus and LMS page.
- ➜ Revisit the policy throughout the semester, particularly before major assignments and assessments.
- ➜ Reflect on the policy with students at the end of the semester; treat the policy as an ongoing conversation.

3.2 Shaping the policy with your students

Instructors who want more student buy-in can involve students directly in parts of the policy-setting process. Teaching-center guidance from institutions including UC Berkeley, Columbia, UW–Madison, Virginia Commonwealth, and MIT recommends involving students in ongoing discussions about course AI-use policies (Appendix B collects these references). NYU models the principle at two scales: the university's student-authored AI-use guidelines (New York University, 2026), written by NYU students for NYU students with contributors credited by name, and the Provost's faculty workshop "From Policing to Building Trust" (Duff, 2026), which recommends co-creating the course AI-use policy with students, with the instructor keeping the final call.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

3.1 Regardless of which AI policy you adopt

3.2 Shaping the policy with your students

Case study 1

Case study 2

3.3 A student-facing checklist and copy-ready templates

Suggested assignment box

Suggested student AI disclosure statement

3.4 Examples: writing clear AI instructions

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Ideally, the instructor brings a starter policy in which the fixed, non-negotiable items are explained up front (those tied to accreditation, safety, or core learning objectives, for instance) alongside a small number of open decisions (which disclosure format to use, or where exactly the line falls for limited support on a given assignment type) that students discuss and vote on. The instructor makes the final call and can decline any option that would undermine the course's learning objectives.

This approach has now run in two settings, independently. In two lecture-based courses on responsible AI, students debated five open decisions and voted; in a project-based graduate capstone, taught by one of this guide's working-group leads, students shaped the policy through structured anonymous input. Both used the same basic structure: a starter policy with fixed items and open Red/Yellow/Green decisions. The questions and answers differed across the two settings, reflecting the purpose of involving students in the process.

Case study 1: Shaping AI policy in two Responsible AI courses

The Preface describes the pilot: two technical courses on responsible AI, a required undergraduate course of about 190 students and a graduate elective of about 35, ran the same policy exercise working from a draft of this guide. This box shows the instrument itself.

Each course published a starter policy with three parts: the staff commitments on their own AI use; fixed items, each with a stated rationale (quizzes and the final exam AI-free and device-free; disclosure required; students responsible for everything they submit); and amendable defaults for labs and the course project. Five decisions were left open, worded identically in both courses:

1. **Homework, code.** Yellow (AI may explain and critique your code, but not write it) or Green (AI may write code, with disclosure).
2. **Homework, essays.** Red (no AI), Yellow (think first: your own draft, with AI at the edges), or Green (free use, with disclosure).
3. **Assigned readings.** Red (engagement first: read and think before asking AI), Yellow (free use, with AI-free accountability where it counts), or Green (free use, with disclosure).
4. **Disclosure format.** A brief general statement or an itemized log, and whether "no AI was used" must be stated explicitly.
5. **Verifying understanding.** Random spot-check conversations about submitted work, explanations only for cause, or explain-your-work components built into assignments.

Discussion ran asynchronously on the course forum during week one, facilitated by teaching assistants, with a substantive post expected from each student; participation carried a small amount of course credit (1% of the grade, equivalent to one lab), and was nearly universal. Posts were anonymous to peers but visible to staff. Voting followed via a form. Where lab sections agreed, the majority position was adopted; where they split, the instructor decided and published the reasoning. The adopted policies went into both syllabi in week two, with a mid-semester amendment window on the calendar.

Two results stand out. Students were often stricter than the staff expected: on essays, majorities in both courses (60% of undergraduates, 64% of graduate students) chose Yellow over Green. And the cohorts resolved identical questions differently: undergraduates wanted explanations of their work only for cause (82%), while graduate students preferred random spot-checks (68%); undergraduates chose Green for homework code (84%), while 39% of graduate students preferred Yellow.

The activity and voting-form templates are linked in the References section; adapting them to another course takes about 30 minutes.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

3.1 Regardless of which AI policy you adopt

3.2 Shaping the policy with your students

Case study 1

Case study 2

3.3 A student-facing checklist and copy-ready templates

Suggested assignment box

Suggested student AI disclosure statement

3.4 Examples: writing clear AI instructions

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Case study 2: Shaping AI policy in a graduate project-based experiential learning course

The capstone course serves about 70 full- and part-time graduate students of varied ages, academic backgrounds, and countries of origin, working in teams for outside sponsors. In Week 1 the instructor introduced a baseline AI-use policy already drafted into the syllabus; in Week 2 the class began to socialize it, responding anonymously to structured questions over the following week. The design premise, drawn from the instructor's practice in community science and participatory governance, was that policy should be shaped with the people expected to operate under it, not introduced only when someone violates it, and that students should reason first about learning, cognitive responsibility, professional judgment, and trust rather than penalties and detection. Participation was voluntary, uncompensated, and discussed under Chatham House principles; questions averaged a 29% response rate.

Two of the five open questions matched Case Study 1 (**disclosure format; verifying understanding)**. The other three reflected team- and client-based work: **written team deliverables** (Red: no AI in writing; Yellow: AI as a writing aid; Green: transparent AI-assisted writing), **professional communications** (Red: no AI; Yellow: AI as a communication aid; Green: transparent AI-assisted communication), and **feedback and conflict management** (Red: no AI in reflection or feedback; Yellow: AI as a preparation aid; Green: AI-assisted preparation).

Students viewed AI as a useful academic tool when used intentionally, transparently, and within clear boundaries; they supported AI assistance while keeping responsibility for learning, reasoning, and verification with the student; and they favored accountability over surveillance, accepting verification and explain-your-work expectations while resisting burdensome monitoring. For this course, students held that context determines appropriate use: purpose, extent, team expectations, sponsor requirements, and the nature of the task set the boundary.

The exercise also surfaced an open question Case Study 1 could not: when does one member's AI use become the team's responsibility? Students wanted clearer escalation and accountability rules for team settings, including whether a team should face consequences for use it did not know about or approve. This remains an open question, an addition to the Section 4.6 list for any project-based course, and the next subject for student input here.

Across both pilots, students engaged seriously, preferred accountability to surveillance, and treated context, what the work is for, as the test of appropriate use. The differences reflected the individual course contexts: which decisions were worth opening, and where each community drew its lines.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

3.1 Regardless of which AI policy you adopt

3.2 Shaping the policy with your students

Case study 1

Case study 2

3.3 A student-facing checklist and copy-ready templates

Suggested assignment box

Suggested student AI disclosure statement

3.4 Examples: writing clear AI instructions

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

3.3 A student-facing checklist and copy-ready templates

Students may encounter different AI-use policies across courses, and even across assignments in the same course. Instructions should make clear whether AI can be used, how, what cannot be done with it, why the boundary exists, and whether disclosure is required. Before sharing an AI-use policy or assignment instruction, check whether a student can answer:

- Can I use AI?
- What can I use it for?
- What can I not use it for?
- Why are these boundaries in place?
- Which tools or types of tools does the policy apply to?
- Do I need to disclose my use? If so, how?
- What should I do if my intended use isn't addressed?

Our position on disclosure parallels our position on AI detectors: we recommend it, but not as an enforcement tool. Its value is pedagogical. A disclosure requirement makes AI use something students reflect on and can discuss openly, and it gives the instructor a record that can start a conversation, not a confession to audit. Expect imperfect compliance, keep disclosure low-stakes to encourage honest reporting, and never treat a disclosure, or its absence, as evidence of misconduct on its own.

Suggested Assignment Box

> *AI Use for This Assignment — Status: Green (AI-Integrated) / Yellow (Limited AI) / Red (AI Prohibited). You may use AI for: _____. You may not use AI for: _____. Disclosure required: Yes / No. If yes, submit: _____. Why: _____. Questions? Ask the instructor before submitting.*

Suggested student AI disclosure statement

> *"AI Use Statement: I used [tool] to support [brainstorming / outlining / debugging / editing / another activity]. I used the output to [brief explanation]. I reviewed and modified the output before incorporating it into my work and remain responsible for the final submission." An instructor may also offer a "No AI Use Statement" option: "No generative AI tools were used in preparing this assignment."*

Where formal attribution matters, the disclosure statement can become a citation: APA, MLA, and Chicago each provide citation formats for AI-generated content, typically an in-text citation naming the model and date (e.g., OpenAI, 2024), comparable to citing a quoted source, and in papers the tools used can be named in the methods or acknowledgments section. Where verification matters, ask students to retain the prompt and the model name, version, and date, so use can be documented when required.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

3.1 Regardless of which AI policy you adopt

3.2 Shaping the policy with your students

Case study 1

Case study 2

3.3 A student-facing checklist and copy-ready templates

Suggested assignment box

Suggested student AI disclosure statement

3.4 Examples: writing clear AI instructions

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

3.4 Examples: writing clear AI instructions

The table below presents five recurring cases, each with typical unclear wording and a clearer version to adapt. The clearer versions all do the same three things: name the permitted and prohibited uses, give the reason, and state the disclosure expectation.

Scenario	Unclear, and why it fails	Clearer
Red (AI Prohibited)	"The use of AI is prohibited." No reason given; unclear whether grammar checkers, autocomplete, translation tools, or AI-enabled search count.	"Generative AI tools may not be used to complete this assignment, including tools such as ChatGPT, Gemini, or Claude to generate answers, explanations, code, outlines, or written content. Independent problem-solving is one of the learning objectives being assessed, and AI-generated assistance would interfere with that assessment. Standard spell-check and basic grammar correction are permitted. If you're unsure whether a tool is allowed, ask before using it."
Yellow (Limited AI)	"AI can be used for brainstorming but not for completing the assignment." Where does brainstorming end? Is an outline brainstorming? A rewritten paragraph?	"You may use generative AI during the early stages of your work to generate ideas, identify possible research questions, create a preliminary outline, or explain concepts. You may not use AI to write the final response, generate the final analysis, or produce text you submit as your own. This assignment assesses your ability to develop and communicate the final argument independently. If you use AI during the permitted stages, include a short acknowledgment describing how."
Green (AI-Integrated)	"Students may use AI tools for this assignment." "Use" is undefined: brainstorming, writing, editing, generating code, or producing the entire assignment could all qualify.	"You may use generative AI tools for brainstorming, exploring alternative approaches, debugging code, and receiving feedback on your work, including to help explain concepts you don't understand. The final analysis, decisions, and submitted work must reflect your own understanding, and you are responsible for verifying the accuracy of any AI-generated information you use. At the end of your submission, include a brief statement identifying the AI tool(s) you used and how."

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

3.1 Regardless of which AI policy you adopt

3.2 Shaping the policy with your students

Case study 1

Case study 2

3.3 A student-facing checklist and copy-ready templates

Suggested assignment box

Suggested student AI disclosure statement

3.4 Examples: writing clear AI instructions

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Scenario	Unclear, and why it fails	Clearer
Course baseline vs. assignment	"AI is permitted in this course unless otherwise stated." A student may miss the exception, or assume the course policy covers every assignment.	"The default policy for this course allows the use of generative AI as a learning and support tool. Individual assignments may set different AI expectations depending on the skill being assessed; each major assignment will state its own AI-use policy, which overrides the course default." Then use the assignment box from Section 3.3 on each assignment.
Disclosure	"AI use must be cited." What information, where, and is a full chat transcript required?	"If you use generative AI for this assignment, add a short AI Use Statement at the end of your submission identifying the tool, what you used it for, and how you reviewed or modified its output. You do not need to submit your complete chat transcript unless specifically requested."

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

4.1 Why AI-related misuse is not plagiarism

4.2 Relation to the existing Codes of Conduct

4.3 Instructor guidance: generative AI and academic integrity

Setting expectations in the syllabus

If you become aware of a potential violation

What counts as evidence

Why not rely on an AI detector?

4.4 Proposed wording updates to a Code of Conduct

4.5 What peer institutions do

A few representative examples (as of summer 2026)

4.6 Open decisions above the course level

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

4. Academic-Integrity Process

This section covers what happens if a student is suspected of AI misuse: why AI misuse does not map cleanly onto existing misconduct categories, how the standard academic-integrity process works, plain-language guidance for handling a suspected case, how Codes of Conduct can be updated, how peer institutions approach the problem, and the decisions each institution still has to make.

4.1 Why AI-related misuse is not plagiarism

Unauthorized AI use is not the same problem as plagiarism, on four counts:

1. **No "wronged author."** Plagiarism has a victim whose credit was taken; AI model output does not necessarily have an individually identifiable source whose credit was taken. The harm instead falls on the reader or evaluator who is deceived about provenance. Existing Code of Conduct guidelines remain in place in instances where AI paraphrases identifiable or copyrighted sources.

2. **Repair centers on disclosure, not attribution.** Disclosing AI use tells the evaluator about the process by which the text was made, not where it was copied from.

3. **Responsibility for output and process differs**. Writers may delegate some control over the writing process to AI, but not responsibility for accuracy. This differs from plagiarism, where they are presenting another person's work as their own.

4. **The evidence base is different.** Plagiarism can be shown by matching a source document; AI detectors are comparatively unreliable (high false-positive rates), and may put multilingual students at particular risk of false accusation.

Two decisions run through the remainder of this section: the default when a course or assignment is silent, and whether disclosure is independently required, so that undeclared use is itself a violation even where the assistance was permitted. We collect these, with the other open decisions, in Section 4.6.

4.2 Relation to the existing Codes of Conduct

Typical Codes of Conduct at colleges and universities already cover cheating, fabrication, plagiarism, unauthorized collaboration, duplicating work, and forgery, and set out a multi-step process. There are variations, but the general process is as follows:

1. **Instructor stage:** if a student is suspected of academic dishonesty, the instructor collects evidence and discusses the allegations with the student, offering them an opportunity to respond.

2. **Resolution or referral:** if the student admits the violation, the instructor imposes a sanction within their authority or refers the case to the department/school/administration. If the student denies the violation or disputes the sanction, the matter goes to the administration and, if unresolved, to a committee for student affairs.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

4.1 Why AI-related misuse is not plagiarism

4.2 Relation to the existing Codes of Conduct

4.3 Instructor guidance: generative AI and academic integrity

Setting expectations in the syllabus

If you become aware of a potential violation

What counts as evidence

Why not rely on an AI detector?

4.4 Proposed wording updates to a Code of Conduct

4.5 What peer institutions do

A few representative examples (as of summer 2026)

4.6 Open decisions above the course level

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

3. **Hearing:** the case is reviewed by a panel of faculty and/or students; typically a violation is found on a preponderance-of-the-evidence standard (more likely than not).

4. **Sanctions:** sanctions range from rejection of the assignment (with a substitute) to a failing grade on the assignment or course, to dismissal, suspension, or expulsion, depending on severity.

5. **Appeal:** sanctions can be appealed.

This layering is not hypothetical at NYU, where this guide originated. The university-wide Academic Integrity policy sets the common definitions of plagiarism and cheating, and expressly provides that each school may establish "its own detailed supplemental guidelines for academic integrity, consistent with its own culture" (New York University, n.d.-a). The updates proposed in Section 4.4 and Appendix E operate at exactly that supplemental layer, within the NYU Tandon School of Engineering's Student Code of Conduct (NYU Tandon School of Engineering, n.d.).

4.3 Instructor guidance: generative AI and academic integrity

Institutions should give instructors plain-language guidance for suspected cases, independent of whether Code of Conduct changes are adopted; what follows reflects our experience and conversations. Section 3 covers introducing your policy to students; the expectations below additionally protect both the student and the instructor if a suspected case arises.

Setting expectations in the syllabus

- ➜ Require students to verify the accuracy of their citations, references, and factual claims. It is an important skill whether or not AI was involved.
- ➜ Prepare students to explain their methods and sources on request and in person; their response may determine whether a formal referral is needed.
- ➜ When stating which tools, methods, or references are permitted for an assignment (Section 3.3), note also what deviating from those limits may indicate.
- ➜ Tell students in advance that they should expect to explain submitted work verbally. This statement supports any oral-defense step within a Code of Conduct investigation process.

Instructors should state the policy twice: first in the syllabus and verbally in class, then again before the first major assignment, since students may still be adding and dropping courses throughout the first few days or weeks of the academic term.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

4.1 Why AI-related misuse is not plagiarism

4.2 Relation to the existing Codes of Conduct

4.3 Instructor guidance: generative AI and academic integrity

Setting expectations in the syllabus

If you become aware of a potential violation

What counts as evidence

Why not rely on an AI detector?

4.4 Proposed wording updates to a Code of Conduct

4.5 What peer institutions do

A few representative examples (as of summer 2026)

4.6 Open decisions above the course level

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

If you become aware of a potential violation

Treat an initial impression that a submission may involve unauthorized AI use as the start of an inquiry, not a finding. An informal, in-person conversation with the student is strongly encouraged before any formal referral, since it may resolve or narrow the concern considerably. Where the conversation raises further concern, document it and proceed formally by the usual processes.

How the inquiry proceeds depends on the course's default AI policy. Where AI use is prohibited by default, first confirm what AI uses were explicitly permitted before determining whether the student violated the policy. Where AI use is limited or integrated by default, first confirm what AI uses were explicitly prohibited before determining whether the student violated the policy. In either case, the conversation with the student (not a detector score, which is too unreliable) should carry the most weight.

The University of Liverpool (2024) provides a useful distinction, treating cases as either poor academic practice (honest misunderstanding of the rules, handled informally with the student, no referral) or academic misconduct (deliberate deception, handled through the formal process described in Section 4.2); most first cases are the former.

What counts as evidence

- The work cites or relies on sources, data, or quotations that do not exist or cannot be verified (fabrication).
- The work includes methods or reasoning that exceed the scope of the course, or seem implausible without unauthorized assistance.
- The student cannot verbally explain, expand on, or defend the submitted work when asked.
- The work is inconsistent with the student's earlier drafts, outlines, or office-hours discussion of the same assignment.

In cases of academic dishonesty, findings usually rest on a preponderance of the evidence (more likely than not), which still requires more than an instructor's suspicion or a single detector reading. Multiple independent forms of evidence provide stronger support for a finding.

Why not rely on an AI detector?

Detection of generative AI outputs is a moving target: reliability shifts with every new model, and at the time of writing, detectors misclassify human-written text as AI-generated in roughly 10–20% of cases, with substantially higher false-positive rates for multilingual and non-native English writers (Liang et al., 2023). A detector score, flag, or indicator alone therefore does not meet the preponderance standard; it should be paired with the direct, verifiable evidence and student conversation described above. This is also NYU's advice to its own faculty: the Provost's workshop series recommends against AI-detection tools entirely, on grounds of accuracy, bias, and student privacy, noting that uploading student work to a third-party AI tool without consent can violate FERPA (Duff, 2026). Building trusting relationships with students, and designing assessments that make unauthorized AI use less useful, will likely do more to protect academic integrity than any detection tool available today.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

4.1 Why AI-related misuse is not plagiarism

4.2 Relation to the existing Codes of Conduct

4.3 Instructor guidance: generative AI and academic integrity

Setting expectations in the syllabus

If you become aware of a potential violation

What counts as evidence

Why not rely on an AI detector?

4.4 Proposed wording updates to a Code of Conduct

4.5 What peer institutions do

A few representative examples (as of summer 2026)

4.6 Open decisions above the course level

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

4.4 Proposed wording updates to a Code of Conduct

Most conduct codes predate generative AI, and updating one is an exercise in rewording, not rewriting: the definitions and evidence rules need to name AI explicitly, and the rest of the machinery can stand. We ran this exercise on NYU Tandon School of Engineering's Code of Conduct, clause by clause; the full worked review is in Appendix E, as an illustration of the exercise, not as language to copy verbatim. The two examples below show the two moves that recur throughout it.

Pre-AI Text (Example)	Suggested Reword	Why Change Is Needed
"Academic misconduct is defined here as misrepresentation, deception, dishonesty, or falsification committed by a student to influence a grade or other academic evaluation."	"Academic misconduct is defined here as misrepresentation, deception, dishonesty, or falsification, **including the unauthorized or undisclosed use of generative artificial intelligence or other automated tools to produce, complete, or substantially assist in academic work** by a student to influence a grade or other academic evaluation."	Definitions should name AI rather than leave it to inference. The phrase "unauthorized or undisclosed use" treats disclosure as a separate requirement, independent of any level of authorization. Other institutions may choose more permissive stances.
"Evidence of academic dishonesty or misconduct must be collected and presented to the student within a reasonable time frame."	"Evidence of academic dishonesty or misconduct **(which may include metadata, version history, or draft materials, and which must be weighed alongside instructor judgment rather than relied upon in isolation)** must be collected and presented to the student within a reasonable time frame."	AI-detection tools have a high false-positive rate and cannot carry an accusation on their own; naming this in the code protects both students and instructors.

Every edit in the full review does one of these two things: it names AI explicitly where a definition assumed all work is human-made (misconduct, cheating, fabrication, collaboration, and one new category, unauthorized content generation), or it governs how AI-related evidence and sanctions are handled. The full set reflects one possible stance, a restrictive default with a broad disclosure requirement; the underlying decisions are collected in Section 4.6.

Reviewing a Code of Conduct through an AI-policy lens rarely stays contained to the Code of Conduct alone. The same pre-AI gaps tend to surface anywhere else your institution defines misconduct: in department-level honor codes, TA and grading guidelines, thesis and dissertation integrity policies, and research-conduct policies. Treat Appendix E as a model for that broader review, rather than a one-time edit to a single document: once your core Code of Conduct is updated, it is worth reviewing other official documents that reference or restate its definitions.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

4.1 Why AI-related misuse is not plagiarism

4.2 Relation to the existing Codes of Conduct

4.3 Instructor guidance: generative AI and academic integrity

Setting expectations in the syllabus

If you become aware of a potential violation

What counts as evidence

Why not rely on an AI detector?

4.4 Proposed wording updates to a Code of Conduct

4.5 What peer institutions do

A few representative examples (as of summer 2026)

4.6 Open decisions above the course level

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

4.5 What peer institutions do

Many U.S. research universities have developed new approaches to generative AI and academic integrity. Four patterns recur, and each bears directly on the open decisions in Section 4.6:

- **A default matters.** Most high-ranked private institutions (Harvard, Stanford, Princeton, Columbia) treat unauthorized AI use as unpermitted assistance when a syllabus is silent; some public institutions and programs (e.g. the University of Washington's Foster School of Business) flip the default, allowing AI unless the instructor forbids it. Very few leave the default undefined.
- **Decentralization is the norm.** MIT, Penn, Michigan, Duke, Penn State, and others delegate the policy decision to instructors or departments rather than mandating one campus-wide rule, consistent with the light-touch preference documented in Section 1.
- **AI detectors are widely rejected as primary disciplinary evidence.** Stanford, MIT, Caltech, Yale, Michigan, and Duke discourage or ban detection software (Turnitin, GPTZero) as primary evidence, citing false-positive rates and documented bias against non-native English writers.
- **Equity and access are a live concern.** Banning AI without providing a secure institutional alternative risks pushing students toward unvetted public tools and unequal access to premium models; MIT, Harvard, Michigan, and Johns Hopkins provide such alternatives ("U-M GPT," "HopGPT").

A few representative examples (as of summer 2026)

Institution	Default When Syllabus Is Silent	Notable Feature
Stanford University	Restrictive: treated analogously to assistance from another person.	Does not support campus-wide AI-detection software; flags detectors as unreliable and biased.
Princeton University	Restrictive: AI output cannot be cited as a source; use must be disclosed if permitted.	Ended its 133-year unproctored exam tradition, all in-class exams are proctored from July 1, 2026 (Bonette, 2026).
University of Washington (Foster School)	Permissive: AI use allowed without disclosure by default unless the instructor states otherwise.	Students must still be able to verbally defend and explain any submitted work.
Johns Hopkins (Whiting School of Engineering)	Mandatory labeling: every course must display Red/Yellow/Green.	Suspected violations cannot be resolved informally; must go to a formal Academic Ethics Board Panel.
University of Texas at Austin	Mandatory labeling: instructors choose AI Permitted / Prohibited / Partially Permitted.	Explicitly prohibits faculty and TAs from using AI to grade or assess student work.

These five examples are illustrations, not a survey: they span restrictive to permissive defaults, and purely decentralized policy-making to mandatory syllabus labeling. Benchmarking against a regional or peer-institution consortium can provide a fuller picture reflecting your own constraints.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

4.1 Why AI-related misuse is not plagiarism

4.2 Relation to the existing Codes of Conduct

4.3 Instructor guidance: generative AI and academic integrity

Setting expectations in the syllabus

If you become aware of a potential violation

What counts as evidence

Why not rely on an AI detector?

4.4 Proposed wording updates to a Code of Conduct

4.5 What peer institutions do

A few representative examples (as of summer 2026)

4.6 Open decisions above the course level

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

4.6 Open decisions above the course level

Whichever stance a school or department takes, five decisions need explicit answers; peer practice (Section 4.5) shows workable versions of each.

1. **The default when a course or assignment is silent.** Should AI assistance be prohibited by default absent instructor permission, or permitted by default absent instructor prohibition? Peer practice splits on this (Section 4.5), and a school-level default should apply only when no more specific course or assignment rule has been given.
2. **Whether a disclosure requirement applies independently of that default,** so that undeclared use is itself a violation, enforceable even where the underlying assistance was permitted.
3. **Consistent language on AI detection across the whole document set:** no template, syllabus statement, or conduct code should suggest that a detector score by itself establishes a violation; enforcement language should match the evidence standard in Section 4.3.
4. **The purpose of disclosure, stated consistently:** disclosure records how AI was used. It does not by itself show that a student mastered the material, and it does not by itself prove misconduct (see Section 1.3, Practice 3).
5. **Whether to provide a short default policy for silent syllabi** that instructors can point to even if they have not written their own.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

A note on how this guide was prepared

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Conclusion

This guide gives structure to one of the hardest questions facing instructors today: how much AI use is compatible with what a specific assignment is supposed to teach. That answer is different for a first-year problem set, a capstone design project, and a reflective essay, and no menu can make the decision for you. This guide offers a shared vocabulary, tested policy language, socialization practices, and a fair process so that each decision is easier to make, easier to state clearly, easier to explain to students, and easier to handle consistently when something goes wrong. Our advice fits in one sentence: set a baseline and tune it per assignment, expect a profile rather than a single color, shape the policy with your students, and verify mastery AI-free.

Instructors should treat their own policy as a starting point rather than a finished product. At NYU Tandon, where this guide originated, the working group's next step is to track how these policies actually perform once departments adopt them. A systematic review will look at outcomes like the gap between take-home and AI-free assessment, the volume and disposition of academic-integrity cases, and instructor and student experience with the Red (AI Prohibited) / Yellow (Limited AI) / Green (AI-Integrated) model and the socialization practices. We encourage any institution adapting this guide to build in a similar review from the start, rather than treating a policy, once written, as settled. The landscape scan established where we started; the planned review will show what happens after the policies are put into practice.

If you adapt this guide for your own department or institution, tell us what worked, what did not, and what you added. The questions in Section 4.6 are still open, and experience across institutions can help inform them.

A note on how this guide was prepared

Literature and peer-institution research throughout this guide were located with the help of AI search and synthesis tools (ChatGPT-5.6 Luna, Claude Sonnet 5, Mistral Medium 3.5), then reviewed and selected by working-group members. Figure 1 was generated with Claude Opus 5.5, using the quick user guide as an initial prompt. All policy language, analysis, and editorial judgment are the working group's own.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Institutional policies and guidance cited in this guide

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

References

Bonette, J. (2026, May 12). After 133 years, Princeton is going back to proctoring exams. Princeton Alumni Weekly. https://paw.princeton.edu/article/after-133-years-princeton-going-back-proctoring-exams

Choudhuri, R., Sanchez, C. A., Burnett, M., & Sarma, A. (2026). Thinking less, trusting more: GenAI's impacts on students' cognitive habits [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2601.22430

Duff, D. L. (2026, January). From policing to building trust: Reframing academic integrity in the age of AI [Workshop]. Generative AI Workshops, New York University Office of the Provost. https://nyu.edu/ai

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. British Journal of Educational Technology, 56(2), 489–530. https://doi.org/10.1111/bjet.13544

Furze, L., Perkins, M., Roe, J., & MacVaugh, J. (2024). The AI Assessment Scale (AIAS) in action: A pilot implementation of GenAI-supported assessment. Australasian Journal of Educational Technology, 40(4), 38–55. https://doi.org/10.14742/ajet.9434

Gonsalves, C. (2025). Addressing student non-compliance in AI use declarations: Implications for academic integrity and assessment in higher education. Assessment & Evaluation in Higher Education, 50(4), 592–606. https://doi.org/10.1080/02602938.2024.2415654

Kofinas, A. K., Tsay, C. H.-H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. British Journal of Educational Technology, 56(6), 2522–2549. https://doi.org/10.1111/bjet.13585

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2506.08872

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7), Article 100779. https://doi.org/10.1016/j.patter.2023.100779

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. Journal of University Teaching and Learning Practice, 21(6). https://doi.org/10.53761/q3azde36

Perkins, M., Roe, J., & Furze, L. (2025). How (not) to use the AI Assessment Scale. Journal of Applied Learning & Teaching, 8(2), 14–23. https://doi.org/10.37074/jalt.2025.8.2.15

Rojas-Galeano, S. (2025). New kid in the classroom: Exploring student perceptions of AI coding assistants [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2507.22900

Schilke, O., & Reimann, M. (2025). The transparency dilemma: How AI disclosure erodes trust. Organizational Behavior and Human Decision Processes, 188, Article 104405. https://doi.org/10.1016/j.obhdp.2025.104405

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Institutional policies and guidance cited in this guide

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Stoyanovich, J. (2026). AI-use policy voting form template [Online form template]. NYU Center for Responsible AI. https://r-ai.co/ai-use-voting-template

Strathern, M. (1997). "Improving ratings": Audit in the British university system. European Review, 5(3), 305–321. https://doi.org/10.1002/(SICI)1234-981X(199707)5:3%3C305::AID-EURO184%3E3.0.CO;2-4

Strömberg, D., Lei, V., & Wu, Y. (2026). The generative AI learning penalty: Evidence from Chinese secondary education (CEPR Discussion Paper No. 21577). CEPR Press. https://cepr.org/publications/dp21577

Institutional policies and guidance cited in this guide

California Institute of Technology. (n.d.). AI in teaching. https://aiinteaching.caltech.edu/

California Institute of Technology, Information Management Systems and Services. (n.d.). Generative AI. https://www.imss.caltech.edu/services/ai

Carnegie Mellon University, Eberly Center for Teaching Excellence & Educational Innovation. (n.d.). Examples of possible academic integrity policies that address student use of generative AI tools. https://www.cmu.edu/teaching/technology/aitools/academicintegrity/index.html

Columbia Business School, Office of Student Affairs. (n.d.). Generative AI policy. https://students.business.columbia.edu/office-of-student-affairs/academic-advising-and-student-success/academic-integrity/generative-ai-policy

Columbia University, Center for Teaching and Learning. (n.d.). Developing AI syllabus statements and communicating with students about AI. https://ai.ctl.columbia.edu/explore/ai-syllabus-statement/

Columbia University, Office of the Provost. (n.d.). Generative AI policy. https://provost.columbia.edu/content/office-senior-vice-provost/ai-policy

Cornell University, Center for Teaching Innovation. (n.d.). AI course policy icons. https://teaching.cornell.edu/generative-artificial-intelligence/ai-course-policy-icons

Duke University, Center for Teaching and Learning. (n.d.). Artificial intelligence policies: Guidelines and considerations. https://ctl.duke.edu/ai-and-teaching-at-duke-2/artificial-intelligence-policies-in-syllabi-guidelines-and-considerations/

Harvard University, Faculty of Arts and Sciences, Office of Undergraduate Education. (n.d.). Generative AI guidance. https://oue.fas.harvard.edu/faculty-resources/generative-ai-guidance

Harvard University Information Technology. (n.d.). Generative AI guidelines. https://www.huit.harvard.edu/ai/guidelines

Johns Hopkins University, Whiting School of Engineering, Engineering for Professionals. (n.d.). Student guidance on responsible generative AI use. https://ep.jhu.edu/student-services/academic-services/academic-policies/student-guidance-on-responsible-generative-ai-use/

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Institutional policies and guidance cited in this guide

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Massachusetts Institute of Technology, Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training. (2026). AI and education. https://aiandeducation.mit.edu/

New York University. (n.d.-a). Academic integrity for students at NYU. https://www.nyu.edu/about/policies-guidelines-compliance/policies-and-guidelines/academic-integrity-for-students-at-nyu.html

New York University. (n.d.-b). Frequently asked questions about teaching and AI. https://www.nyu.edu/faculty/teaching-and-learning-resources/teaching-with-generative-tools/frequently-asked-questions.html

New York University, Department of Computer Science. (2026). CS202-001: Policies and grading [Course policy, M. Walfish, Fall 2026]. https://cs.nyu.edu/~mwalfish/classes/26fa/policies.html

New York University, Department of Computer Science. (n.d.). CSCI-UA 102: Data structures [Course syllabus, J. Klukowska]. https://cims.nyu.edu/~joannakl/cs102/syllabus.html

New York University, Department of Mathematics. (n.d.). Academic integrity. https://math.nyu.edu/dynamic/undergrad/ba-cas/academic-integrity/

New York University. (2026). Student learning with Generative AI. https://www.nyu.edu/faculty/teaching-and-learning-resources/Student-Learning-with-Generative-AI.html

NYU Tandon School of Engineering. (n.d.). Student Code of Conduct. https://engineering.nyu.edu/life-tandon/student-life/student-advocacy/student-code-conduct

Pennsylvania State University, Academic Integrity Office. (n.d.). Artificial intelligence. https://integrity.psu.edu/ai

Princeton University. (n.d.). Disclosing generative AI. Scholarly Integrity. https://scholarlyintegrity.princeton.edu/students/disclosing-generative-ai-use

Stanford University, Center for Teaching and Learning. (n.d.). Governing artificial intelligence [Example course policy, INTLPOL 245B / LAW 4052]. https://ctl.stanford.edu/opportunities/governing-artificial-intelligence

Stanford University, Graduate School of Business, Teaching and Learning Hub. (n.d.). Course policies on generative AI use. https://tlhub.stanford.edu/docs/course-policies-on-generative-ai-use/

Stanford University, Office of Community Standards. (n.d.). Generative AI policy guidance. https://communitystandards.stanford.edu/generative-ai-policy-guidance

Stanford University, Program in Writing and Rhetoric. (n.d.). PWR policy on the use of GenAI and LLMs. https://pwr.stanford.edu/about-pwr/pwr-policies/pwr-policy-use-genai-and-llms

Teachers College, Columbia University, Digital Futures Institute. (n.d.). TC guidance. https://www.tc.columbia.edu/digitalfuturesinstitute/ai-in-education/artificial-intelligence-in-education/tc-guidance/

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Institutional policies and guidance cited in this guide

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

University College London. (n.d.). Three categories of GenAI use in assessment. https://www.ucl.ac.uk/teaching-learning/artificial-intelligence-education/using-ai-tools-assessment

University of California, Berkeley, Center for Teaching & Learning. (n.d.). Communicating course policies and talking with students. https://teaching.berkeley.edu/teaching-strategies/ai-teaching-learning/communicating-course-policies-and-talking-students

University of Liverpool. (2024). Code of practice on assessment, Appendix L: Academic integrity policy 2024–25. https://www.liverpool.ac.uk/media/livacuk/tqsd/code-of-practice-on-assessment/2024-25,appendix_L_cop_assess.pdf

University of Michigan. (n.d.). Course policies & syllabi statements. U-M Generative AI. https://genai.umich.edu/resources/faculty/course-policies

University of Pennsylvania, Center for Excellence in Teaching, Learning, and Innovation. (n.d.). Creating AI policies. https://cetli.upenn.edu/resources/generative-ai/course-policies-communication/

University of Texas at Austin, Office of the Executive Vice President and Provost. (n.d.). Your syllabus at UT Austin. https://provost.utexas.edu/the-office/academic-affairs/your-syllabus-at-ut-austin/

University of Washington, Foster School of Business. (n.d.). AI syllabus statement. https://foster.uw.edu/about-foster-school/ai-at-foster/ai-syllabus-statement/

University of Wisconsin–Madison, Center for Teaching, Learning & Mentoring. (n.d.). AI statements for course syllabi. https://ctlm.wisc.edu/generativeai/ai-statements-for-course-syllabi/

Virginia Commonwealth University, Center for Teaching and Learning Excellence. (n.d.). Gen AI guidelines for students. https://faculty.provost.vcu.edu/faculty-resources/center-for-teaching-and-learning-excellence/gen-ai-guidelines-for-students/

Yale University, Office of the Provost. (2023, September 20). Guidelines for the use of generative AI tools. https://provost.yale.edu/news/guidelines-use-generative-ai-tools

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Policy examples reviewed for this guide

If you want to fully allow / encourage AI use

If you want to allow some AI use

If you want to prohibit AI use

Appendix B.

Appendix C.

Appendix D.

Appendix E.

Appendix A.

Policy examples reviewed for this guide

This appendix lists the course, program, and school AI-use policies and the teaching-center template collections, behind the summary in Section 2.2. Entries are grouped by the level of AI use they illustrate: Red (AI Prohibited), Yellow (Limited AI), or Green (AI-Integrated). Some template collections span multiple levels and are included where relevant. As Section 2.2 explains, we cite these as evidence of the range of current practice, not as models to adopt. Links were verified in summer 2026; these policies are revised frequently and may have changed since then.

If you want to fully allow / encourage AI use

1. Carnegie Mellon Eberly Center: Six copy-ready models spanning prohibition, full encouragement, assignment-specific permission, and ungraded classroom use. https://www.cmu.edu/teaching/technology/aitools/academicintegrity/index.html
2. NYU CS Data Structures: uses AI as a tutor, feedback source, and programming partner; students document, verify, and explain the work, while exams remain AI-free. https://cims.nyu.edu/~joannakl/cs102/syllabus.html
3. Stanford Governing AI course: broad permission paired with disclosure and full student responsibility for the final work. https://ctl.stanford.edu/opportunities/governing-artificial-intelligence
4. UCL: a simple three-category assessment model: no AI, assistive AI, or AI-integrated work. https://www.ucl.ac.uk/teaching-learning/artificial-intelligence-education/using-ai-tools-assessment

If you want to allow some AI use

1. Cornell AI Course Policy Icons: assignment-level labels for permitted tools, attribution, privacy, exploration-only use, and AI-free work. https://teaching.cornell.edu/generative-artificial-intelligence/ai-course-policy-icons
2. Harvard FAS: restrictive, mixed, and fully encouraging templates; the mixed version disallows AI use unless an assignment says otherwise. https://oue.fas.harvard.edu/faculty-resources/generative-ai-guidance
3. NYU Courant Mathematics: allows students to check work with tools like ChatGPT, but not to copy or submit answers without their own reasoning. https://math.nyu.edu/dynamic/undergrad/ba-cas/academic-integrity/
4. Teachers College, Columbia: adaptable templates for assignment-specific use, purpose-specific use, prior approval, and use with disclosure. https://www.tc.columbia.edu/digitalfuturesinstitute/ai-in-education/artificial-intelligence-in-education/tc-guidance/
5. UW–Madison: three syllabus options plus a detailed AI-use statement covering tool, version, purpose, prompt, and retained records. https://ctlm.wisc.edu/self-serve-resources/generativeai/ai-statements-for-course-syllabi

If you want to prohibit AI use

1. Stanford PWR: protects writing skills by prohibiting AI composition or revision of major assignments, while allowing instructors to authorize limited support. https://pwr.stanford.edu/about-pwr/pwr-policies/pwr-policy-use-genai-and-llms
2. Columbia Business School: treats AI as unauthorized unless the instructor gives permission and requires an explanation of permitted use. https://students.business.columbia.edu/office-of-student-affairs/academic-advising-and-student-success/academic-integrity/generative-ai-policy
3. Harvard FAS: also offers a maximally restrictive template forbidding AI at every stage of the work process. https://oue.fas.harvard.edu/faculty-resources/generative-ai-guidance
4. NYU Courant CS strict course policy: prohibits generative AI at every stage because the course is designed to develop independent systems thinking and programming skill. https://cs.nyu.edu/~mwalfish/classes/26fa/policies.html

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.
Institutional guidance on communicating and shaping AI policy with students

Appendix C.

Appendix D.

Appendix E.

Appendix B.

Institutional guidance on communicating and shaping AI policy with students

The guidance below, from teaching centers and university initiatives, is the basis for the recommendations in Sections 3.1 and 3.2; as in Appendix A, we cite it as evidence of converging practice. Links verified in summer 2026.

NYU (university-wide): NYU's student-authored AI-use guidelines, written by NYU students for NYU students with contributors credited by name, model the same principle at university scale: the people a policy governs help to write it. The Provost's faculty workshop "From Policing to Building Trust" (Duff, 2026) recommends the same practice: co-create the course AI-use policy with students, with the instructor keeping the final call. https://nyu.edu/ai

UC Berkeley: *"A clearly stated course policy on GenAI use helps ensure that students understand not only what is allowed, but also why. Policies are more effective when they are grounded in pedagogical reasoning and tied to the specific learning goals of the course … a good GenAI policy should not merely act as a set of rules but should invite students into a shared conversation."* https://teaching.berkeley.edu/teaching-strategies/ai-teaching-learning/communicating-course-policies-and-talking-students

Columbia University: *"Since instructors and academic departments have the ability to develop their own AI course policies, students may encounter different guidelines across their courses. To avoid confusion, we encourage all instructors to include a clear, well-defined generative AI statement in their course syllabi that clearly delineates whether and how students can use AI in their coursework."* https://ai.ctl.columbia.edu/explore/ai-syllabus-statement/

UW-Madison: Instructors are encouraged to share expectations both at the start of a course and throughout the semester, via syllabi, the LMS, and course discussions, to promote integrity and prevent confusion. https://ctlm.wisc.edu/generativeai/ai-statements-for-course-syllabi/

Virginia Commonwealth: Guidance frames the goal as "not to set hard do's and don'ts, but to provide guidance and have conversations about AI use," to be updated as the technology changes. https://faculty.provost.vcu.edu/faculty-resources/center-for-teaching-and-learning-excellence/gen-ai-guidelines-for-students/

MIT: Policies should state their rationale and be tied to learning objectives; no one-size-fits-all; detectors are not evidence. https://aiandeducation.mit.edu/

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.
Components of a worked course profile

Appendix D.

Appendix E.

Appendix C.

Components of a worked course profile

This appendix walks through building a course profile, the exercise Sections 2.2 and 2.3 describe, using a single worked example. First, set the course baseline and justify it; then write one entry per assignment or assignment type, overriding the baseline as needed.

Set the baseline. This is the course default; assignments can override it.

- ☐ **Red (AI Prohibited)**
  Students complete the work without generative AI, because independent performance is part of what the course teaches.
- ☐ **Yellow (Limited AI)**
  AI may be used only for listed purposes (explaining concepts, checking work), not to produce what is submitted.
- ☐ **Green (AI-Integrated)**
  AI use is expected and disclosed; students verify outputs and remain responsible for the final submission.

Provide a justification for the baseline.

➔ **Justification:**

In the worked examples that follow, the baseline is Yellow: most of the work in the course is analysis that students must be able to perform and explain themselves. The four entries then show how assignments can override the default in either direction and in how much detail an assignment-specific policy can provide.

Homework 1 and the Labs stay Yellow but give specific justifications and concretely explain how AI may or may not be used, and whether and how to provide disclosure. The project moves up to Green, because professional audit work is done with AI in the loop; the exam moves down to Red, because it certifies unaided mastery. Together they illustrate the core recommendation in miniature: the course has a profile, not a single color. Individual course profiles will differ.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.
Components of a worked course profile

Appendix D.

Appendix E.

Homework: Yellow

- **Justification:** This assignment asks you to audit a classifier for performance disparities across groups; computing the fairness metrics and interpreting what they do and don't tell you is the skill being assessed, and you build it by doing the analysis yourself.
- **AI may be used to:** explain concepts, check solutions after an attempt
- **AI may not be used to:** solve assigned problems, write submitted answers
- **AI-free check:** explain one problem during the lab
- **Disclosure:** state in which problems you used AI and how, for example: *"Problem 2: after computing the false positive rate gap myself, I asked ChatGPT to check my reasoning about why the metrics disagree; the interpretation I submitted is my own."*

Labs: Yellow

- **Justification:** Labs are practice for the homework. They are graded on completion, and we are not looking for polished submissions, so this is your opportunity to learn the concepts.
- **AI may be used to:** explain concepts, help debug your code, explore different options after your own attempt
- **AI may not be used to:** complete the lab exercises for you
- **AI-free check:** none
- **Disclosure:** not required

Project: Green

- **Justification:** The project asks you to produce an audit of a real system, and in professional practice that work is done with AI tools in the loop; using them well, verifying what they produce, and standing behind the result is part of the skill being assessed.
- **AI may be used to:** help write and debug code, draft and revise text, explore analyses and alternatives, summarize related work
- **AI may not be used to:** fabricate results or citations, produce analysis you have not verified and cannot explain
- **AI-free check:** the project presentation; every team member should be prepared to explain and defend any part of the work in Q&A
- **Disclosure**: an AI Use section in the project report: which tools were used, for what, and how you verified the outputs

Final exam: Red

- **Justification:** The exam certifies that you can reason about the core concepts of this course on your own.
- **AI may be used to:** prepare for the exam, but not on the exam itself
- **AI may not be used to:** complete any part of the exam; the exam is in class, pen and paper, no devices
- **AI-free check:** the exam is the check
- **Disclosure:** not applicable

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

AI-Use Policy: Starter Draft for Deliberation

Thread D1: Coding assignments

Thread D2: Written reports

Thread D3: Readings

Thread D4: Disclosure

Thread D5: Explaining your work

Thread D6: General (amendments and anything else)

Appendix E.

Appendix D.

AI-Use Policy: Starter Draft for Deliberation

This appendix reproduces the starter policy posted in the undergraduate pilot course at the start of Fall 2026, lightly edited for length; it is the instrument behind Case Study 1 in Section 3.2. Students read this post in week one, argued their positions in the six discussion threads over several days, and then voted; the adopted policy went into the syllabus. It is reproduced in its original student-facing voice on purpose, because the tone is part of the method: the staff's commitments come first, every fixed rule carries its reason, and the open decisions state their tradeoffs honestly. One vocabulary note: the post predates this guide's final terminology, and its "unlimited transparent use" options correspond to the guide's Green (AI-Integrated). The discussion and voting form templates for running this exercise in your own course are linked in the References.

What course staff commit to, regardless of what we adopt:

- We will not use AI detectors as evidence of anything. They are unreliable and biased against non-native English writers, and this course will not pretend otherwise.
- If we ever suspect a policy violation, the process starts with a conversation with you, not with a detector score.
- We won't use AI to grade your work or write feedback on it.
- We will never upload your work to generative AI platforms without your awareness and consent.
- Where course materials are prepared with AI assistance, we will say so.

Fixed components of the policy

- **Quizzes (in lab) and the final exam are AI-free and device-free.** This is where you show us what you know on your own, and because we know these are your own work, we can be flexible everywhere else. Here is why this matters: in one Tandon engineering course, homework scores averaged close to 100% while proctored exams on the same material averaged about 50%. We don't want that to be you.
- **Disclosure is required whenever AI is used on graded work.** Undisclosed AI use is considered cheating and results in an F in the course. What the disclosure must contain is up for decision (D4).
- **Do not upload course materials,** including but not limited to assignment prompts, solution code, quiz or exam content, to AI tools without explicit permission from the instructor.
- **You remain responsible for everything you submit:** that it is correct and that it is your work, within whatever AI rules we adopt.

Defaults you can propose to change (use the General thread):

- **Labs:** AI allowed as a learning aid (explain a concept, unblock you, help debug) after you have attempted each step yourself. Labs are graded on completion, so there is nothing to gain from using AI to make the notebook look better; use it to understand the material.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.
AI-Use Policy: Starter Draft for Deliberation

Thread D1: Coding assignments

Thread D2: Written reports

Thread D3: Readings

Thread D4: Disclosure

Thread D5: Explaining your work

Thread D6: General (amendments and anything else)

Appendix E.

- **Team project:** transparent use. AI may support brainstorming, coding, analysis, and writing; the team reviews, verifies, and substantially revises anything AI-generated, discloses tools and their contributions, and any member may be asked to explain any part of the submission.

Thread D1: Coding assignments

Your three homeworks are mostly programming with data. The learning objective at stake: implementing and understanding the techniques yourself (fairness interventions, explanations, audits), because in later weeks and in the project you will need them as working knowledge, not as something you once had generated for you. Select one option.

- **Option A (Yellow), some AI in coding.** AI may explain concepts, help you debug your own code, and check your solution after an independent attempt. AI may not generate code you submit. Disclose which problems you used it on.
- **Option B (Green), unlimited transparent use of AI in coding.** AI may help write submitted code. You must test and verify it, be able to explain and reproduce any part without AI, and disclose which parts AI contributed to.

The honest tradeoff: Option A protects practice (early studies of AI-assisted programming learners find they feel more confident but do worse when asked to extend their own code unaided). Option B matches how many of you will work professionally and treats you as the responsible engineer of record. Post your position and your reason.

Thread D2: Written reports

Each homework includes a one-to-two page written report interpreting your results: what the data mean, for whom, and with what caveats. The learning objective at stake: this is where the actual thinking of responsible data science happens, and writing is how that thinking forms. It is also the skill employers and graduate programs will assume you have. Select one option.

- **Option A (Red), no AI in writing.** Spelling and grammar tools are fine; no generative AI in the writing at any stage.
- **Option B (Yellow), some AI in writing.** You may use AI to brainstorm before drafting and to get feedback on your finished draft. The sentences you submit are written by you.
- **Option C (Green), unlimited transparent use of AI in writing.** AI may help draft; you must substantially revise, verify every claim, disclose its role, and be able to defend every sentence as your own position.

Thread D3: Readings

The course has weekly readings, and the lectures and quizzes assume you have engaged with them deeply, not just read a summary. The learning objective at stake: the readings help you understand the ideas in depth, and the quizzes (which are AI-free) test your understanding. Select one option.

- **Option A (Red), no AI in reading.** You read the assigned material yourself, in full, with no AI summaries, explanations, or notes standing in for it. What you bring to lecture and quizzes is your own reading.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

AI-Use Policy: Starter Draft for Deliberation

Thread D1: Coding assignments

Thread D2: Written reports

Thread D3: Readings

Thread D4: Disclosure

Thread D5: Explaining your work

Thread D6: General (amendments and anything else)

Appendix E.

- **Option B (Yellow), some AI in reading.** You read the assigned material yourself; after reading, you may use AI to clarify a concept, check your understanding, or explore connections. AI summaries do not replace the reading.
- **Option C (Green), unlimited transparent use of AI in reading.** You prepare however you like, including with AI summaries before or instead of the full reading. You are accountable for depth of understanding (or lack thereof) and you disclose AI use as appropriate.

Worth knowing: MIT's committee put "summarizing assigned readings in place of engaging directly with the original material" on its list of AI uses that most directly deprive students of learning.

Thread D4: Disclosure

Disclosure is required (fixed). The question is what it must contain to be useful. Select one option.

- **Option A, minimal disclosure.** One or two sentences per assignment: tool, purpose, extent, and what you did to verify. Example: "I used Claude to debug my preprocessing code in Q2 and to suggest structure for the essay outline. I tested all the code, and the analysis and conclusions are my own."
- **Option B, detailed disclosure.** Identify the specific questions, sections, or code blocks AI influenced, and keep your prompts available if asked. More burden for the student, more information for the grader.

Thread D5: Explaining your work

A fixed component of this policy is that you are responsible for the work you submit. The question is how this can be checked.

- **Option A, for-cause explanations only.** The explain-your-work conversation may happen when staff have a specific concern.
- **Option B, spot-check and for-cause explanations.** In addition to for-cause explanations, course staff may ask any student to walk through any submitted work in office hours; inability to explain your own submission is treated as a problem in itself. Applied occasionally and at random, announced as normal practice, not accusation.

Thread D6: General (amendments and anything else)

Propose amendments to the defaults (labs, team project), or to anything in the starter policy, including the fixed items. Fixed items need a strong argument to be changed, but strong arguments are exactly what this course is about. This is also the place for language, access, and accommodation considerations: if a rule or a disclosure requirement would impact multilingual writers or students with approved accommodations differently, say so here. And anything else that the five threads don't cover.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.
A Worked Code of Conduct Review

Appendix E.

A Worked Code of Conduct Review

This appendix contains the full clause-by-clause review of NYU Tandon School of Engineering's Code of Conduct, summarized in Section 4.4. Highlighted text marks the proposed insertions. As with the other appendices, this is a worked illustration of the exercise, not language to copy verbatim. The edits reflect one possible stance: a restrictive default with a broad disclosure requirement. Section 4.6 collects the underlying decisions, and the exact wording will depend on your own institution's existing conduct code and educational mission. We would expect the specific edits to look different elsewhere even where the underlying gaps are the same.

Pre-AI Text (example)	Suggested Reword	Why Change Is Needed
It is through the process of submitting their own work and receiving honest feedback on that work that students may progress academically. Those who breach the School's rules on academic integrity may be subject to sanctions, up to and including dismissal, suspension, or expulsion from the University under this Code.	It is through the process of submitting their own **original work, produced within the bounds of any generative AI use permitted for that assignment,** and receiving honest feedback on that work that students may progress academically. Those who breach the School's rules on academic integrity, **including rules governing the use of generative AI or other automated tools,** may be subject to sanctions, up to and including dismissal, suspension, or expulsion from the University under this Code.	The introduction assumes that originality of 'own work' is self-evident, but Gen AI blurs the line between student-authored and LLM-authored content.
Academic misconduct is defined here as misrepresentation, deception, dishonesty, or falsification committed by a student to influence a grade or other academic evaluation.	Academic misconduct is defined here as misrepresentation, deception, dishonesty, or falsification, **including the unauthorized or undisclosed use of generative artificial intelligence or other automated tools to produce, complete, or substantially assist in academic work** by a student to influence a grade or other academic evaluation.	Definitions should name AI rather than leave it to inference. The particular phrase "unauthorized or undisclosed use" indicates that we view disclosure as a separate requirement, independent of any level of authorization. Other institutions may choose more permissive stances depending on their needs.
Students should review the class syllabi for more details to obtain a clear understanding of any further specific rules or standards of academic misconduct.	Students should review the class syllabi for more details to obtain a clear understanding of any further specific rules or standards of academic misconduct, including policies surrounding the use of generative AI. **In the absence of specific course policies stated in the syllabus, the use of generative AI tools (including chatbots, agents, browser extensions, or AI-enabled search features or devices) for completing coursework (in whole, in part, or as part of the process) is not permitted. In cases where use of such tools is permitted, it must also be disclosed. Finally, course materials may not be uploaded to AI tools without express permission from the instructor.**	A default policy is needed so that student and faculty expectations are clear. The example language here is fairly restrictive, but leaves open the possibility of students using Gen AI for other purposes e.g. clarifying concepts, general Q&A on the material, searching for relevant books/documents/media, etc. Again, the particulars of the policy will be necessarily institution-dependent.

Preface

Overview: Origin and How to Use This Guide

1. Landscape & Scoping Analysis

2. Building Your Course Policy

3. Student-Facing Socialization

4. Academic-Integrity Process

Conclusion

References

Appendix A.

Appendix B.

Appendix C.

Appendix D.

Appendix E.

A Worked Code of Conduct Review

Cheating is defined as intentionally using or attempting to use unauthorized notes, books, electronic media, or communications in an exam; talking with fellow students or looking at another person's work during an exam; submitting work prepared in advance for an in-class examination; having someone take an exam for you or taking an exam for someone else; violating other rules governing the administration of examinations.	Cheating is defined as intentionally using or attempting to use unauthorized notes, books, electronic media, **devices, communications, or generative AI tools (including chatbots, browser extensions, or AI-enabled search features or devices)** in an exam; talking with fellow students or looking at another person's work during an exam; submitting work prepared in advance, **including AI-generated drafts,** for an in-class examination; having someone **or something, including an AI tool or AI-enabled device,** take an exam for you or taking an exam for someone else; violating other rules governing the administration of examinations.	Electronic media and communications pre-date Gen AI, which should now be mentioned explicitly, but still in general terms to avoid stale or dated policies in a rapidly changing environment. We also recommend explicitly adding language around electronic and/or AI-enabled devices to include, for example, smart glasses.
Fabrication includes but is not limited to, knowingly falsifying experimental data and/or citations.	Fabrication includes, but is not limited to, knowingly falsifying experimental data and/or citations, **or submitting data, citations, quotations, or sources generated or altered by an AI tool without verifying their accuracy and existence.**	Generative AI can produce fabricated data, citations, quotations, or sources that students and faculty may not recognize as such.
Unauthorized collaboration is defined as working together on work that was meant to be done individually.	Unauthorized collaboration is defined as working together on work that was meant to be done individually, **including with a generative AI tool used as a substitute collaborator or co-author, absent express permission from the course instructor.**	"Working together" presumes another person so students might think collaboration with a chatbot doesn't count. Institutions should clarify their stance. Again, this is a more restrictive posture.
	Unauthorized content generation is the production of academic work, in whole or in part, for academic credit, progression, or award, whether or not payment or another favor is involved, using unapproved or undeclared human or technological assistance.	We recommend introducing a new category of academic misconduct since plagiarism does not directly map onto misuse of Gen AI tools. This is one such example.
Evidence of academic dishonesty or misconduct must be collected and presented to the student within a reasonable time frame.	Evidence of academic dishonesty or misconduct, **(which may include metadata, version history, or draft materials, and which must be weighed alongside instructor judgment rather than relied upon in isolation),** must be collected and presented to the student within a reasonable time frame.	AI-detection tools have a high false-positive rate, so they can't be relied upon on their own. Faculty should take this into consideration when bringing forward accusations.
Mild sanctions (e.g. assignment resubmission) may be particularly appropriate where there is an honest misunderstanding about the degree of collaboration permitted by a course instructor, or where a course instructor is not sure a student has intentionally acted dishonestly.	Mild sanctions (e.g. assignment resubmission) may be particularly appropriate where there is an honest **misunderstanding about the degree of collaboration or the extent of generative AI use permitted by a course instructor,** or where a course instructor is not sure a student has intentionally acted dishonestly.	AI permission will vary within courses; naming it in CoCs, especially with respect to sanctions, can help to alleviate confusion over AI integration and rules.