"AI assists. You think. You analyze. You write. You take responsibility": Creating a Course AI Use Policy Template --Policy Text, Justification and Rule Summary for My Law & Religion Class at Penn State Dickinson

Larry Catá Backer (白 轲)
(in collaboration with Harvey AI; Gemini, Claude, Grok, ChatGPT)
W. Richard and Mary Eshelman Faculty Scholar; Professor of Law and International Affairs
Penn State Dickinson Law, a unit of The Pennsylvania State University system | 239 Lewis Katz Building, University Park, PA
16802 1.814.863.3640 (direct) || lcb11@psu.edu

Abstract: This essay provides the text of a Model Classroom AI Class Use Policy along with a justification for its terms. It is grounded on the premise that a course AI use policy should be built around a justified allocation of risk rather than a simple permit/prohibit binary. The author reasons that categorical bans are unenforceable and misaligned with how AI is already embedded in legal practice, while unrestricted use would erode the ability to assess unaided student analysis. The resulting design principle is a narrow, conditional exception layered onto a default of independent work, with permission calibrated to where oversight and iterative disclosure are feasible.

That reasoning drives each structural choice. The exception is confined to the final paper (not exams or problem sets) because papers unfold over time, allowing disclosure and instructor engagement that timed assessments cannot accommodate. Permitted and prohibited uses are framed illustratively, not as a bright-line test, because forcing research, analysis, and writing into mechanically verifiable categories would create false confidence at the margins; ambiguity is instead deployed deliberately to discourage boundary-engineering. By contrast, incorporated AI-generated content is subject to fixed numerical ceilings (100 words per instance, 10% aggregate), because that narrower, mechanical act of quotation is better suited to a bright-line rule than the broader assistance-versus-substitution question. Mandatory naming, a full prompt-and-response appendix, in-text attribution, and a signed certification convert disclosure into an enforceable condition of the exception itself, so partial compliance affords no defense, and the institution retains audit authority independent of prior suspicion.

The author frames the resulting template as a first iteration, refined through human review and testing against several AI systems (including Harvey AI), to be revised further based on classroom experience.

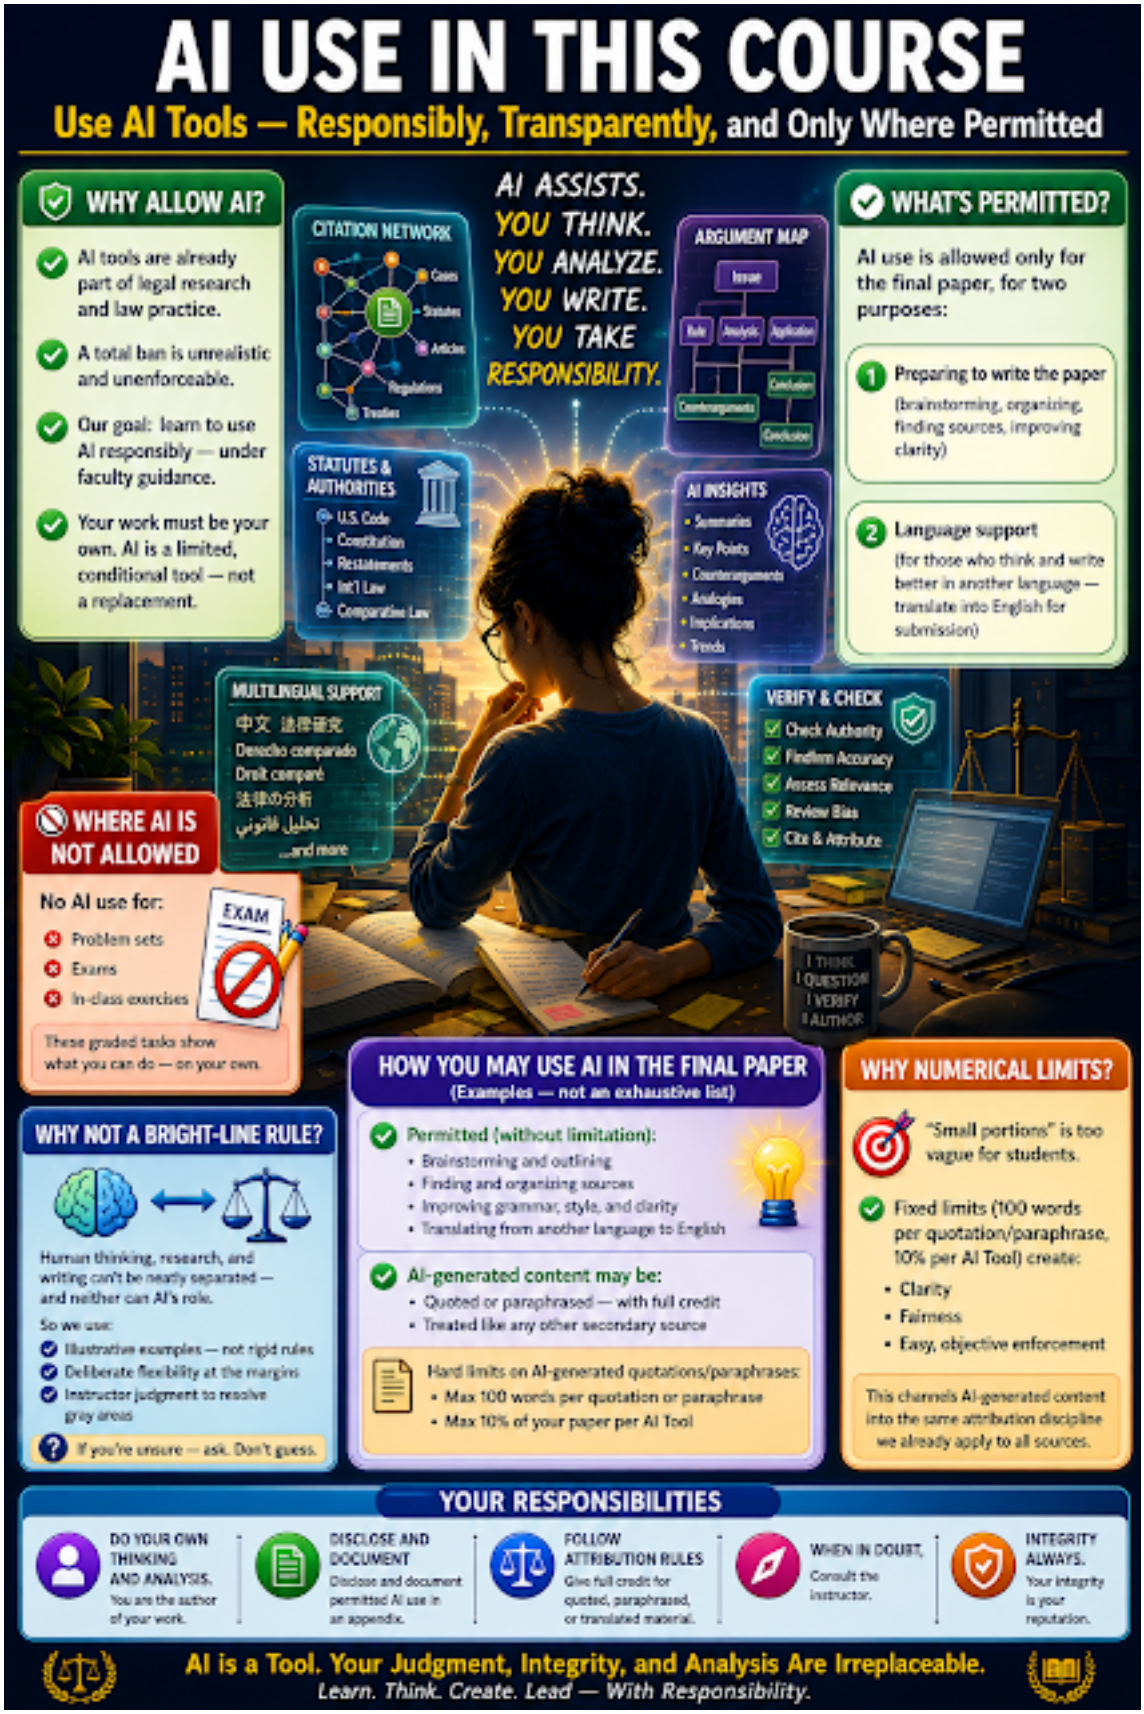
AI USE IN THIS COURSE
Use AI Tools — Responsibly, Transparently, and Only Where Permitted
WHY ALLOW AI?
AI tools are already part of legal research and law practice.
A total ban is unrealistic and unenforceable.
Our goal: learn to use AI responsibly — under faculty guidance.
Your work must be your own. AI is a limited, conditional tool — not a replacement.
CITATION NETWORK
Cases
Statutes
Articles
Regulations
Treaties
AI ASSISTS.
YOU THINK.
YOU ANALYZE.
YOU WRITE.
YOU TAKE
RESPONSIBILITY.
ARGUMENT MAP
Issue
Rule
Analysis
Application
Counterarguments
Conclusion
Conclusion
STATUTES & AUTHORITIES
U.S. Code
Constitution
Restatements
Int'l Law
Comparative Law
AI INSIGHTS
Summaries
Key Points
Counterarguments
Analogies
Implications
Trends
WHAT'S PERMITTED?
AI use is allowed only for the final paper, for two purposes:
1 Preparing to write the paper
(brainstorming, organizing, finding sources, improving clarity)
2 Language support
(for those who think and write better in another language — translate into English for submission)
MULTILINGUAL SUPPORT
中文 法律研究
Derecho comparado
Droit comparé
法律の分析
تحليل قانوني
...and more
VERIFY & CHECK
Check Authority
Confirm Accuracy
Assess Relevance
Review Bias
Cite & Attribute
WHERE AI IS NOT ALLOWED
No AI use for:
Problem sets
Exams
In-class exercises
EXAM
These graded tasks show what you can do — on your own.
I THINK
I QUESTION
I VERIFY
I AUTHOR
HOW YOU MAY USE AI IN THE FINAL PAPER
(Examples — not an exhaustive list)
Permitted (without limitation):
Brainstorming and outlining
Finding and organizing sources
Improving grammar, style, and clarity
Translating from another language to English
AI-generated content may be:
Quoted or paraphrased — with full credit
Treated like any other secondary source
Hard limits on AI-generated quotations/paraphrases:
Max 100 words per quotation or paraphrase
Max 10% of your paper per AI Tool
WHY NOT A BRIGHT-LINE RULE?
Human thinking, research, and writing can't be neatly separated — and neither can AI's role.
So we use:
Illustrative examples — not rigid rules
Deliberate flexibility at the margins
Instructor judgment to resolve gray areas
If you're unsure — ask. Don't guess.
WHY NUMERICAL LIMITS?
"Small portions" is too vague for students.
Fixed limits (100 words per quotation/paraphrase, 10% per AI Tool) create:
Clarity
Fairness
Easy, objective enforcement
This channels AI-generated content into the same attribution discipline we already apply to all sources.
YOUR RESPONSIBILITIES
DO YOUR OWN THINKING AND ANALYSIS.
You are the author of your work.
DISCLOSE AND DOCUMENT
Disclose and document permitted AI use in an appendix.
FOLLOW ATTRIBUTION RULES
Give full credit for quoted, paraphrased, or translated material.
WHEN IN DOUBT,
Consult the instructor.
INTEGRITY ALWAYS.
Your integrity is your reputation.
AI is a Tool. Your Judgment, Integrity, and Analysis Are Irreplaceable.
Learn. Think. Create. Lead — With Responsibility.

"AI assists. You think. You analyze. You write. You take responsibility."

(Course AI Use Model Slogan suggested by ChatGPT)

I have been writing about the challenges of figuring out if, whether or how to incorporate or use of AI Tools (generally including large language models, neural networks, and other computational, generative, or agentic systems) by students (faculty have their own problems) in coursework.

I started by looking at the way the U.S. legal academy has, to date, sought to respond to the challenge, as it is euphemistically labelled from out of the linguistic word salad jungle that is contemporary American bureaucratic languages.[1] This report consolidates a multi-stage analysis of generative artificial intelligence (AI) coursework and examination policies across a retrieval-based sample of twelve American law schools and programs. The study demonstrates that law school AI governance cannot be reduced to a single linear spectrum; instead, policies vary independently along three distinct structural axes: default polarity (restrictive versus permissive baselines), drafting style (determinate rules versus interpretive standards), and a two-tier autonomy structure (governing institution-to-instructor and instructor-to-student relationships). Cross-analysis reveals that substantive restrictiveness does not predict structural design, meaning schools with identical baselines often impose vastly different interpretive or administrative burdens on students and faculty.

I then shifted gears and asked the machine systems themselves (or at least five of them, all U.S.-centric or at least US/Anglo/European centric--Harvey AI, Claude, Grok, ChatGPT, and Gemini) what they thought (to the extent that we transpose human notions of thinking onto the computational environment in which machine systems operate) taking humans into account.[2] The analysis summarizes each system's reasoning and resulting policy text, compares their structural choices along default polarity, drafting style, and autonomy architecture, and categorizes the five outputs by these dimensions. It gives particular attention to the functional divergence between Gemini's tool-based tiering (classifying software by computational architecture) and ChatGPT's task-based categorization (classifying assignments by information dependence), and assesses the practical feasibility of the versioning and archiving practices several systems propose. I got both a set of far more interesting responses and opened a doorway to re-examining or seeing the perhaps inevitability of transforming contemporary analogue and human constrained notions of

[1] Backer, Larry Catá, Structure, Opacity, and Convergence: A Consolidated Analysis of Law School Generative AI Coursework and Exam Policies Larry Catá Backer (轲) (in collaboration with Harvey AI) (July 12, 2026). Available at SSRN: https://ssrn.com/abstract=

[2] Backer, Larry Catá, Rethinking AI Governance in Legal Education --Five Machines (Grok, Harvey, ChatGPT, Claude, and Gemini), One Question, No Consensus but Five Archetypes (July 14, 2026). Available at SSRN: https://ssrn.com/abstract=

law and legal education. These were grounded in conceits about text, time, and the immutability of data.

The last stage was an examination of the possibility of shifting the operational lens of the policy from a substantially to a substantially machine system orientation.[3] This report examines a three-stage experiment the first part of which analyzed U.S. law school efforts at construction AI education policies were considered and against which, in parts two and three, five AI systems-Harvey AI, Claude, ChatGPT, Grok, and Gemini-were pressed to construct governance policies for AI use in law school coursework, first from a ″human-centric″ computational perspective and then, more radically, ″without regard to... human-centric normative guardrails.″ The central finding, confirmed repeatedly by the systems′ own self-audits, is that none achieved genuine machine-centered derivation independent of human normative content; each produced a technically reformulated restatement of pre-existing human intellectual traditions, a fact several systems conceded directly when challenged.

All of that was fun; some of it was--as Anglo-American discourse tends to prefer when they are "solving" problems--immediately useful. Yet none of this solved my immediate problem--thinking through the best way of incorporating or rejecting the incorporation of AI Tools in my classes, or working through some pragmatic middle ground.

I took the institutional-cultural prodding and applied it seriously in the context of my own circumstances (there is nothing like personalizing a challenge to capture one's interest). I will be teaching a course next semester in which I will assign the drafting and submission of a final paper written around the course themes--the constitutional law of the United States, with a glance at parallel developments elsewhere. I took my earlier study of the U.S. approaches as well as the current context in which such policies can be plausibly developed for a law school course in a U.S. law school. I came up with a draft, which, when tested against human reception proved worthy of further revision (the "standard" process for the evolution of such documents, retaining the decision to choose among alternatives because it was my course and for application in specific context decisions are necessary to move from broader frameworks to an actual operational architecture. Lastly I took the last of the human-human drafts and asked a machine system to review it as a function of the parameters described in my review of the state of AI policy in U.S. law schools. I used Harvey AI because the law school where I am based has entered into a relationship with the controllers of that machine system. That also produced some revision; machine systems, when trained to my liking, do not spare ego. And that was appreciated; and a reminder that one always remains a student of drafting for clarity; that clarity, even within perfectly clear text remains elusive; and that sometimes it is the spaces that ambiguity creates that are the most effective pathways to equitable operationalization.

[3] Backer, Larry Catá, Structure, Legitimacy, and the Limits of Machine-Centered Derivation: An Analysis of Five AI Systems′ Third-Stage Attempts to Construct Machine-Centric Governance Policies for Legal Education (July 26, 2026). Available at SSRN: https://ssrn.com/abstract=

So, with all of that in mind I produced a template form for a Course AI Tools use policy. That template derived from a set of six policy principles that I had been using in past years as a sort of default--no use of AI--developed again in the law school in which I am based.

The template is specific to my needs and desires--to permit the use of AI Tools (as I have sort of defined them above) for two purposes; the first is in getting ready to write the paper and the second is in creating a space where people who write and think better in a language other than English can do so and then translate their work product into English for course submission. The parameters of AI use can be broadened and narrowed to suit the instructor with a change in the operative paragraph of the template.

The summary and justification for the text of the course policy template follows here; the text of the policy template follows below. Comments, reactions, etc. welcome offline. This is very much a work in progress that will develop further as data becomes available (how it works with students). And my thanks to the humans and machine systems that helped make this better than it started out.

* * *

Rationale and Justification for the Policy

Why permit AI use at all, rather than prohibit it outright. Generative AI tools are already embedded in legal research platforms, word processors, and the broader practice of law that students will enter. A categorical ban would be difficult to enforce, would not reflect how students will actually work as practicing attorneys, and would forgo the opportunity to teach students how to use these tools responsibly while they are still subject to close faculty oversight. The policy's default in Section 1 nonetheless remains that work must be the student's own; Section 2 is a deliberately narrow, conditional departure from that default, not a general license. And it is constructed so that an instructor can change its parameters to meet needs and objectives, which likely change with each course and over time.

It is fair to point out that if a categorical ban is difficult to enforce, the disclosure-and-appendix regime built to police the narrower exception faces some version of the same verification difficulty; prompt histories can in principle be edited, regenerated, or omitted before submission. No verification system is foolproof, and both machine systems themselves and human experience predating them are unsparing in reminding us of that. That objection, however, is best tested rather than argued away, and my class will serve as my own personal laboratory for doing so. The working hypothesis is not that certification and disclosure make evasion impossible, but that they raise the cost of evasion enough, at the margin, to shift the balance of compliance meaningfully in the right direction. Whether that hypothesis holds is a question for the data, not for the drafting table.

Why the exception is limited to the final paper. For me, confining permitted AI use to the final paper, and prohibiting it without exception for problem sets, exams, and in-class exercises, preserves a low-stakes, frequently graded body of student work as an unambiguous baseline for assessing whether a student can perform the relevant legal analysis unaided. The final paper is treated differently because it is typically produced over a longer period (at least theoretically, knowing that student engagement with final papers is realized within broadly different time horizons), with more opportunity for iterative disclosure, appendix documentation, and instructor engagement than a timed exam permits. But again, there is no reason why an instructor cannot broaden the exception from the "no use" baseline principle to something broader; as broad as they like within the constraints of university or unit policy.

It must be conceded that this justification, resting as it does on time horizon and iterative engagement, does not by itself address equity: students with more experience prompting AI tools, or with access to premium tools, could gain a differential advantage precisely because the paper is the one assignment where use is permitted. Disclosure and certification do not eliminate that differential, but they help expose its extent, which is itself a necessary first step; the data gathered from this initial iteration is meant to inform whether and how the model should evolve. It is also worth noting, more constructively, that a visible differential in AI fluency is not only a

problem to be managed but a potential teaching opportunity: made transparent through disclosure, it suggests a pathway toward using those very differences as a subject of classroom discussion about competent and equitable use, rather than only as a risk to be suppressed.

Why the boundary between permitted and prohibited uses is illustrative rather than a bright-line test. To my mind, a rule that tried to draw a precise, mechanically applied line between permissible "assistance" and impermissible "substitution" would need to disentangle research, analysis, and writing into independently verifiable strands. That, I believe, is not a reliable exercise for human cognitive work, and an attempt to force it risks two modes of failure: (1) false confidence by students that anything outside the enumerated examples is automatically safe, and (2) an unintended chilling effect on legitimate uses that happen to fall near an artificial boundary. Section 2(b)(i) and (ii) are accordingly framed as non-exhaustive examples ("without limitation") rather than a complete catalog, preserving a degree of interpretive ambiguity at the margin. That ambiguity is a deliberate design choice: it is intended to discourage students from engineering borderline uses on the assumption that anything not expressly listed as prohibited is safe. It also shifts interpretation from the student (who generates the interpretive question) to the instructor (who may resolve it in the context of the course, its objectives, and the principles for teaching that specific course).

It is true that shifting interpretation to the instructor after the fact carries a real cost: a student cannot know with certainty, at the moment of use, whether a given use will later be characterized as permitted. Read together with Section 5's "ask first" instruction, though, that same ambiguity does substantial pedagogical work rather than merely imposing an inconvenience. It creates an occasion for direct interaction between instructor and student before the fact, in which the instructor can both assess a proposed use and steer the student toward a better one. That is a teaching moment in its own right, and one that is likely worth the friction that ambiguity introduces, especially where the instructor can use the resulting conversation to engage the student on the broader question of what AI-assisted legal analysis should look like.

Why the quotation/paraphrase proviso exists, and why it now carries hard numerical limits. Section 2(b)(iii) recognizes that AI-generated material can legitimately appear in a paper the same way any other secondary source can: quoted or paraphrased with full authorial credit, rather than silently absorbed into the student's own analysis. This channels any use of AI-generated substantive content into the same attribution discipline students are already expected to apply to human sources, rather than creating a separate, laxer standard for machine-generated content. Unlike the illustrative, non-exhaustive examples in Section 2(b)(i) and (ii), the proviso sets fixed numerical ceilings (a 100-word cap per quotation or paraphrase, and a 10% aggregate cap per AI Tool) rather than relying on an open-ended "small portions" standard.

Bright-line numerical rules are always somewhat problematic for exactly the reason a qualitative standard is not: their precision produces a rigidity that can reduce effectiveness at the margins, where a slightly longer or shorter passage may be substantively indistinguishable from one just inside or outside the line. But that rigidity is the price of a real benefit here. A number disciplines use in a way that a phrase like "small portions" does not; a qualitative standard functions

reasonably well among practitioners who share a professional culture of peer expectations, but for a student it operates as an invitation to bargain over what "small" means, at exactly the moment that incentive is least welcome. The 100-word and 10% figures were not derived empirically; they are meant as a reasonable, if necessarily arbitrary, approximation of what "small portions" was already gesturing toward, chosen because a fixed number is easier for a student to self-apply before submission and easier for an instructor to verify after the fact.

It is also fair to observe, on further reflection, that a per-tool aggregate cap leaves an obvious gap: a student could in principle stay under the 10% ceiling for each of several AI Tools while still incorporating a much larger share of AI-sourced content into the paper overall. To close that gap, an overall aggregate cap of 30% of the paper, measured across all AI Tools combined, has also been added, so that the per-tool ceiling disciplines use of any single tool without leaving multiple-tool use as an unaddressed workaround.

Why disclosure (naming, the appendix, and attribution) is required even for permitted uses. Requiring students to name every AI Tool used, preserve the full sequential prompt-and-response history, and attribute that use in the body of the paper serves three purposes. First, it creates a contemporaneous record that lets an instructor evaluate a specific use against the Section 2(b) standard, rather than relying on the student's own characterization after the fact. Second, it normalizes disclosure as a professional habit, consistent with the emerging expectation that attorneys disclose AI use to clients, courts, and regulators. Third, it shifts the practical burden: a student who complies fully ought to substantially reduce the risk of non-compliance (and where in doubt the template urges students to consult with the instructor before committing themselves to an AI use); a student who used AI in a prohibited way faces a documentary record that makes the underlying conduct difficult to conceal. None of this is airtight, and there will probably be leaks and creative workarounds; the appendix is a record, not a lie detector, and its evidentiary value depends on the certification requirement that follows to give it teeth. But those limits are best tested in the field.

Why the certification requirement exists. The signed certification converts the disclosure obligations into an affirmative representation that can independently support a finding of an academic integrity violation if it proves false, separate from any violation of the underlying use restrictions themselves. This gives the certification independent deterrent value: even a student who correctly predicts that an undisclosed misuse is unlikely to be detected in the appendix still bears the separate risk of having signed a false statement. Treating a false certification as its own violation, rather than folding it into whatever underlying use it misrepresents, is a deliberate escalation, and a warranted one: the certification is the mechanism the entire disclosure structure depends on, and a student's willingness to falsify it goes to candor toward the institution in a way that is analytically distinct from, and arguably more serious than, the underlying use itself. Certification requirements of this kind have, in any event, increasingly become a norm in academic institutions; its use here is neither extraordinary nor burdensome.

Why noncompliance is treated as an unauthorized use rather than a technical deficiency. Section 2(f) states that partial compliance is not a defense, because a disclosure-based system

only works if the disclosure obligations are treated as integral to the permission itself, not as separable paperwork. If a student could use AI Tools in a manner that would otherwise satisfy Section 2(b) but skip the naming, appendix, or certification requirements without losing the exception's protection, the entire enforcement structure would collapse into an honor system with no verification mechanism.

Why the institution retains some review authority regardless of suspicion. Section 2(f) permits review of the appendix and certification for any final paper, not only where a violation is already suspected. An audit-based review structure is intended to function as an ongoing deterrent rather than a one-time gate: students who know their disclosed prompt history may be checked against the submitted paper have an incentive to keep that disclosure accurate and complete throughout the drafting process, not only at the point of submission. All of this, of course, is a function of and constrained by general university and law school policies, including honor codes, academic integrity structures, and the like.

Why this is framed as a first iteration. The template serves as the first step toward its own phenomenology; it may be quite clever conceptually but prove less roadworthy when actually put into operation. This policy is designed to be tested with students in an actual course setting and refined based on what that experience reveals, rather than adopted as a final, fully closed rule. It is, in that sense, a companion piece to my earlier survey of how U.S. law schools have approached generative AI in coursework and exam policies: where that survey described a landscape of considerable structural variation and opacity across institutions, this template is an attempt to draft, in one concrete course setting, a version that is at least internally coherent and testable. The deliberate ambiguity at the margins of Section 2(b), the review authority in Section 2(f), and the certification requirement together function as a first attempt to manage the risk of AI-assisted analysis being passed off as independent student work, without pretending that risk can be eliminated through drafting alone.

Summary Explanation of How the Rules Work

This is what I hope is a more plain-language description of how the course policy template applies to a student preparing work for the course. In this case the language is structured around the assumption that the only exception to the default rule applies to a final course paper. It does not add any new obligations beyond those stated in Sections 1 through 6 above; it is intended only as an explanatory aid.

1. **Default rule.** Unless an instructor says otherwise, all graded or required work must be the student's own. This is the starting point for every other rule in the policy.
2. **Problem sets, exams, and in-class exercises.** AI Tools may not be used to write or modify this work at all, except for the narrow editing aids and research-platform features described in section 4 below, or for translation as described in section 3.
3. **The Work Product.** While this template is drafted with a final paper in mind, the scope of the provision can be any course work product identified by the instructor. A student

may use an AI Tool while preparing the final paper (or otherwise any identified work product), but only if all of the following are true:

- The student's own analysis, arguments, and conclusions drive the paper (or work product). An AI Tool may be used for brainstorming, checking citation form, surfacing counterarguments for the student to evaluate, and grammar or style feedback. An AI Tool may not be used to generate the substantive analysis or prose that the student then adopts with only cosmetic changes.
- If a student wants to include AI-generated material as part of the paper itself (or other work product), it must be handled like a quotation or paraphrase from any other secondary source: credited to the AI Tool by name in a footnote or endnote and in the bibliography, using the citation format described in Section 2(b)(iv), and subject to the 100-word and 10% limits in Section 2(b)(iii). Also included is an overall aggregated cap on use where multiple AI Tools are used of 30% of the paper's word count.
- The student must name every AI Tool used, including its version and how it was accessed.
- The student must attach an appendix to the paper (or other work product as appropriate to the covered work) containing every prompt and response, in the order they occurred.
- The student must include an attribution statement in the body of the paper and sign the certification required by Section 6(b)(i).
- If any of these conditions is not met, the AI use is treated as unauthorized, and it does not matter that the other conditions were satisfied.

4. **Translating work from another language.** A student may use an AI Tool, or a human translator, to translate work into English, but must disclose that translation occurred and how, identify the tool used, make the original-language version available to the instructor on request, and sign the certification required by Section 6(b)(ii). Using a human translator additionally requires the instructor's prior written approval, which will only be granted for a substantial reason. If the translated work is also the final paper, the student must separately comply with the final paper rules in Section 3. As with the final paper (or any other work product that may be described in other variations of Section 2), failing to meet any of these conditions makes the translation-related AI or human-translator use unauthorized, and the instructor or the Law School may review the disclosure, original-language text, and certification for any such work. Use of either AI Tools or a human translator are also subject to certification requirements set out in Section 6.
5. **Grammar checkers and legal research platforms.** Ordinary spelling, grammar, and style-checking tools, and the AI-driven features built into standard legal research platforms (such as case summaries or headnotes), may be used freely and are not subject to the naming, appendix, or certification requirements, as long as the student independently verifies and substantively rewrites any generated text before it goes into the paper rather than copying it in directly.
6. **When in doubt, ask first.** If a student is not sure whether a specific use is permitted, the student should ask the instructor before using the tool, not after submitting the work. And, indeed, the default principle for students is meant to be, if in doubt, ask; if one is tempted to assume, then ask.

7. **What happens if the rules are not followed.** An unauthorized use of an AI Tool may be found to be a violation of academic integrity, and thus subject to the Honor Code's process and, if warranted, its standard sanctions, which can include expulsion or suspension and reporting to bar authorities. The instructor or the Law School may review any student's appendix and certification, whether or not a violation is already suspected, and may compare the disclosed prompts and responses against the paper as submitted. A false certification is treated as a separate violation from whatever underlying use it misrepresents.

ANNEX

Instructor's AI Policy

Instructor: [INSTRUCTOR NAME] **Course:** [COURSE NAME; TERM, YEAR]

Please read the following six sections carefully. We take the Honor Code very seriously at Penn State Law. Violations can result in severe sanctions, including expulsion and suspension, and will be reported to bar authorities. Perhaps the most important idea to take away from the following is that if you are not certain whether using AI technology in a particular way is permissible in a particular course, ask the instructor before using it. The instructor is also always available to address Honor Code questions.

The following constitutes the rules for the use of artificial intelligence, including large language models, neural networks, and other computational, generative, or agentic systems (each, an "AI Tool"), for this class:

1. Expectation of Independent Work. Unless an instructor explicitly provides otherwise, graded or required work submitted by a student for a course or co-curricular activity must be the student's own work. Paragraphs 2 through 6 below describe the specific, instructor-authorized departures from this default that are permitted.

2. Use of Artificial Intelligence. It is a violation of academic integrity for a student to submit work that has been written or modified by another person, or by or through use of any AI Tool as defined above, except as permitted by this Section 2.

(a) Scope of Exception. The exception in this Section 2 applies only to the preparation of the student's final paper for the course. Except as provided in Section 3 (translation) or Section 4 (editing aids and research tools), the prohibition in the preceding paragraph applies without exception to all other graded or required work, including problem sets, exams, and in-class exercises.

(b) Limited and Disclosed Use. The substantive legal analysis, arguments, and conclusions presented in the final paper must originate with the student. An AI Tool may be used only for the purposes described in Section 2(b)(i), and may not be used for the purposes described in Section 2(b)(ii); in all cases, an AI Tool may not be used to generate content that is submitted, without substantive independent analysis and editing by the student, as the student's own work. Subject to the proviso in Section 2(b)(iii) below, in preparing the final paper, a student may use any AI Tool (as defined above) only if all of the following conditions are satisfied:

(i) Permitted uses include, without limitation, brainstorming topics or organizational structures, checking citation form, identifying counterarguments or authorities for the student to independently evaluate, and obtaining feedback on grammar, clarity, or style.

(ii) Prohibited uses include, without limitation, generating the substantive legal analysis, arguments, or conclusions to be adopted with only stylistic changes; drafting sections of prose that are incorporated into the paper without substantive independent rewriting; and generating summaries of legal authority that the student adopts without independently verifying them against the primary source.

(iii) Proviso. Generated content described in Sections 2(b)(i) and (ii) may be used in the final paper only if it is treated as written by someone other than the student, that is it may be directly quoted or paraphrased, only within the following limits:

(A) Length limit. No single quotation or paraphrase of content generated by an AI Tool may exceed 100 words.

(B) Aggregate limit. The total quoted or paraphrased content attributed to any single AI Tool may not exceed 10% of the final paper's total word count. If multiple AI Tools are used total aggregated use of AI Tools cannot exceed 30% of the final paper's total word count.

(C) Exceeding the limits. Use of content generated by an AI Tool beyond the limits in Section 2(b)(iii)(A) or (B) requires the instructor's prior written consent, which will be granted only for substantial cause.

Content used within the limits in this Section 2(b)(iii) must be directly quoted, with the AI Tool given credit as the author, or paraphrased with the same author credit, in a footnote or endnote and in the paper's bibliography or reference section.

(iv) Citation Format. When citing AI-generated work (such as ChatGPT, Gemini, or DALL-E) in a footnote, treat the AI Tool as the author and the developing company as the publisher. Because AI conversations are often non-retrievable, they are cited like personal communications, in the following format: Text/Response to "[Prompt]," [AI Tool and Version], [Company], [Month Day, Year], [URL].

(c) Identification of Tools. The student must identify each AI Tool used, by name, version (if applicable), and the URL or other manner used to access it.

(d) Appendix of Prompts and Responses. The student must submit, as an appendix to the final paper, all prompts submitted to, and all responses generated by, each AI Tool identified under Section 2(c), in the sequential order in which they were generated. Screenshots or downloaded transcripts may be used to satisfy this requirement after consultation with the instructor. The materials identified in this Section 2(d) should be produced in the language in which they were generated; if that language is other than English, an English translation should be provided in accordance with the requirements of Section 3.

(e) Attribution and Certification. The student must comply with the attribution requirement in Section 6 below. In addition, the student must submit with the final paper a signed certification,

in the form set out in Section 6(b)(i), or in a substantially similar form approved by the instructor.

(f) Effect of Noncompliance. Use of an AI Tool that does not satisfy every condition of this Section 2 is unauthorized and constitutes a violation of academic integrity under this Section 2, regardless of whether the student satisfied one or more of the other conditions; partial compliance with this Section 2 is not a defense to a violation of this Section 2, but may be considered in assessing the scope of disciplinary action, subject to the procedures and principles of the Honor Code. The instructor or the Law School may review the appendix submitted under Section 2(d) and the certification submitted under Section 2(e) for any final paper, whether or not there is a specific reason to suspect a violation, including without limitation for assessment and grading, and may compare the prompts and responses disclosed in that appendix against the final paper as submitted.

3. Translation of Work. It is a violation of academic integrity for a student to submit graded or required work first written in another language and then translated into English, whether by another person or by an application, except as permitted by this Section 3.

For the preparation of graded or required work submitted by a student for a course or co-curricular activity, a student may submit work first written in another language and then translated into English by a human translator or by an AI Tool, only if all of the following conditions are satisfied:

(a) Disclosure. The student discloses that an AI Tool was used to translate the work, and describes how it was used. The student must identify each AI Tool used, by name, version (if applicable), and the URL or other manner used to access it.

(b) Submission of the Work in Original Language. At the request of the instructor, the original-language version of any text translated under this Section 3 must be provided to the instructor.

(c) Certification. The student must comply with the attribution requirement in Section 6 below. In addition, the student must submit with the work a signed certification, in the form set out in Section 6(b)(ii), or in a substantially similar form approved by the instructor.

(d) Limitation of Scope of the Translation Exception. This Section 3 applies to work submitted by a student for a course or co-curricular activity. Where such work is also subject to the exception in Section 2 (Use of Artificial Intelligence), the student must comply with the requirements of both Section 2 and this Section 3.

(e) Special Rule for Human Translators. Use of a human translator requires the instructor's prior written approval, which will not be granted absent substantial justification. The certification requirement of Section 6(b)(ii) applies where the instructor gives permission to use a human translator

(f) Effect of Noncompliance. Use of a human translator or an AI Tool that does not satisfy every condition of this Section 3 is unauthorized and constitutes a violation of academic integrity under this Section 3, regardless of whether the student satisfied one or more of the other conditions; partial compliance with this Section 3 is not a defense to a violation of this Section 3, but may be considered in assessing the scope of disciplinary action, subject to the procedures and principles of the Honor Code. The instructor or the Law School may review the disclosure submitted under Section 3(a), the original-language text provided under Section 3(b), and the certification submitted under Section 3(c) for any work subject to this Section 3, whether or not there is a specific reason to suspect a violation, including without limitation for assessment and grading.

4. Editing Aids and Standard Research Tools Permitted. It is not a violation of academic integrity for a student to use an application, including an AI Tool, that identifies grammar, spelling, or stylistic problems and suggests corrections, or to use a legal research platform's built-in features (such as case summaries, headnotes, or search-result ranking) that incorporate an AI Tool, provided that the student does not incorporate that platform's generated text into the final paper without independently verifying it against, and substantively rewriting it from, the underlying primary source. Applications and features used solely as described in this Section 4 are not subject to the conditions in Section 2.

5. When Uncertain, Consult the Instructor. If a student is uncertain whether use of a particular application or AI Tool is permissible in a course or co-curricular activity, the student should consult with the instructor before using it.

6. Attribution Requirement for Instructor-Permitted Use.

(a) General Requirements. Where an instructor explicitly permits the submission of work written or modified by another person or by an application, including an AI Tool (whether under Section 2 or otherwise), the student must provide full attribution and notice of that use in the body of the work, together with any additional information the instructor directs.

(b) Attribution and Certification Pursuant to Sections 2 and 3. In addition to the requirements of Section 6(a), a student submitting graded or required work to which the exception in Section 2 (Use of Artificial Intelligence) or Section 3 (Translation) applies must submit, with that work, a signed certification in the following form, or in a substantially similar form approved by the instructor:

(i) For work to which the exception in Section 2 applies:

"I certify that, except as disclosed in the appendix and attribution statement accompanying this paper, the analysis, arguments, and conclusions presented in this paper are my own. I certify that I have not submitted any content generated by an AI Tool as my own work without substantive independent analysis and editing, and that each AI Tool identified above was used only for the purposes permitted by Section 2(b)(i) and the proviso in Section 2(b)(iii) of this

Policy. I understand that the prompts and responses disclosed in the appendix may be reviewed, that any inconsistency between that appendix and this paper may be treated as evidence of a violation of academic integrity, and that a false certification is itself a separate violation of academic integrity."

(ii) For work to which the exception in Section 3 applies:

[Version where AI Tool was used for translation] *"I certify that, except as otherwise disclosed, this work was translated using an AI Tool, that the use of the AI Tool was limited to translation of the work, and that I have retained a copy of the work in the original language in which it was drafted. I understand that any use of an AI Tool beyond translation requires satisfaction of the requirements of Section 2 of this Instructor's AI Policy, including the further certification required by Section 2(e). I understand that a false certification is itself a separate violation of academic integrity."*

[Version where human translator was used with written instructor approval] *"I certify that, except as otherwise disclosed, this work was translated using the services of a human translator identified below, that the use of the human translator was limited to translation of the work, and that I have retained a copy of the work in the original language in which it was drafted. I understand that any use of a human translator beyond translation is not permitted. I understand that a false certification is itself a separate violation of academic integrity. The name of the human translator and contact details follow ;_____."*