

## Attribution and AI policies plus sample AI use disclosures:

Here is the relevant excerpt from our draft syllabus:

Attribution. You must make sure to provide proper attribution of all sources that you use in your paper. You must read and comply with Section 1(b)(3) of the Law School’s Code of Student Conduct and Responsibility<sup>1</sup> and the Law School Policy on Plagiarism.<sup>2</sup> As the Policy explains, plagiarism is the use of “another’s words or ideas without attribution.” It is important to have a good understanding of how this translates into actual practice. For this purpose, we ask that you read the University of Wisconsin-Madison Writing Center’s handout titled “Acknowledging, Paraphrasing, and Quoting Sources,” which is available at [https://writing.wisc.edu/wp-content/uploads/sites/535/2018/07/Acknowledging\\_Sources.pdf](https://writing.wisc.edu/wp-content/uploads/sites/535/2018/07/Acknowledging_Sources.pdf).

Generative AI. This paragraph contains our policy on your use of generative AI in connection with this seminar. You need not use generative AI in connection with your work for this seminar. If you do use generative AI, you must provide proper attribution (under the standards discussed in the preceding paragraph) for any language that generative AI supplies to you (whether your use of that language consists of a quote or merely a close paraphrase). And, of course, you are responsible for scrutinizing (and ensuring the accuracy of) any ideas or information that generative AI provides to you. Also, if you use generative AI in connection with your assigned work in this seminar (whether that is the two-page paper, your class presentation, the outline of your final research paper, or your final research paper itself), when you hand in (or present) the work in connection with which you used generative AI you must include an “AI use statement” that explains how and when you used the generative AI.<sup>3</sup>

Here are sample AI use disclosures (each of which was generated by ChatGPT o3 in response to prompts that asked it about use cases for specific stages in the course and, with respect to such use cases, asked it to draft a sample AI use disclosure). These are *examples* of some of the ways that I think AI could permissibly be used in our course; I’m giving you not merely the sample use

---

<sup>1</sup> This provision bars “purposeful[], knowing[], or reckless[]” plagiarism “in an academic pursuit[,] including any use of another’s work without attribution, whether such use be verbatim or merely conceptual or structural.” See <https://www.law.upenn.edu/students/policies/conduct-and-responsibility.php>.

<sup>2</sup> The Policy is available here: <https://www.law.upenn.edu/students/policies/conduct-and-responsibility.php>.

<sup>3</sup> Sample AI use statements will be posted separately on our Canvas page.

disclosures but also the preceding dialogue with ChatGPT because that dialogue contains considerations that I'd like you to be aware of if you do use generative AI in connection with our course.

*My prompt: On the readings, I want to stress to the students that they must read all of the required readings themselves. But I suppose that I would be okay with the students using generative AI to summarize the "optional" (as distinct from "required") readings and then using those summaries to determine which of the optional readings to focus on. Would this be a good use? If so, please draft an AI use statement a student could use if they did this for the optional readings on their presentation day.*

### Is it a “good” use?

Yes—provided students (a) **personally read every required reading** and (b) treat the AI-produced optional-reading summaries as a *triage aid* rather than a substitute for engagement. Having ChatGPT (or another model) generate quick abstracts of several optional pieces can help them decide which ones to dive into more deeply for their presentation. The key safeguards are:

1. **Verification.** After choosing an optional article, the student should still skim or read it to confirm that the AI summary is accurate and to pull any quotations or data directly from the source.
2. **Transparency.** Because the AI output directly informed the scope of the assignment, the student must disclose that use in an AI use statement—even if no AI-generated wording ends up in the slides or paper.
3. **No shortcut on required material.** Make clear that summarizing compulsory readings with AI would violate course expectations.

---

## Sample AI Use Statement (for Presentation Day)

### AI use statement for [date] presentation on Smith (1998) and related readings

- **Tool & access date:** OpenAI ChatGPT (model o3), accessed 11 July 2025 via chat.openai.com.
- **Purpose:** To screen five *optional* articles listed for Week 4 and decide which to emphasize in my two-page paper and class presentation.
- **Prompts provided (representative):** “Please give me a concise 200-word summary of ‘Jones, *Law and Markets* (Harvard L. Rev. 2001).’ Highlight its thesis, methods, and main findings.”
- **Outputs used:** I read the resulting summaries, verified key points by skimming each article’s introduction and conclusion, and then selected Jones (2001) and Lee (2015) as the optional works I would integrate.
- **Extent of incorporation:** No AI-generated language appears verbatim or in paraphrase in my submitted paper or slides; the summaries served only to guide my selection of optional readings.
- **Verification steps:** Checked page numbers, quotations, and statistics directly against the original PDFs before final submission.
- **Responsibility:** I remain responsible for the accuracy of all information presented.

This format makes clear **when**, **how**, and **why** the model was used, and reassures the reader that all substantive claims were ultimately traced back to the primary sources.

*My prompt: I don't want the students to use generative AI to actually write anything (or even re-draft anything). But I would be fine with a use like this one: the student first drafts their 2-page paper. Then they upload the paper into generative AI and ask the generative AI to summarize the key points of the paper. If the AI's list of key points isn't the one that they wanted to convey, they re-read the paper and revise it as necessary to make sure their key points are clearly made. What do you think of that use case? And if you think that's a good use case, please draft an AI use statement a student could use if they did this with their draft 2-page paper.*

### **Is this an appropriate use?**

Yes. Treating the model as a **clarity-check mirror** (rather than a ghost-writer) keeps authorship squarely with the student while helping them gauge whether their own prose communicates the intended ideas. Two conditions make it sound:

1. **The draft is student-written.** Nothing from the model flows back into the text except in the student's own words after reflection.
2. **Verification & judgment remain human.** The summary is diagnostic; the student decides what, if anything, to revise.

This is analogous to asking a peer, “What do you hear me saying?”—useful, ethically clean, and unlikely to distort intellectual ownership.

---

### **Sample AI Use Statement (2-Page Paper Clarity Check)**

#### **AI use statement for two-page paper submitted 14 Oct 2025**

- **Tool & date:** Anthropic Claude 3.0 (web interface), 10 Oct 2025.
- **Purpose:** To test whether the draft paper's main arguments were clearly expressed.
- **Prompt (abridged):**  
“Here is my complete draft (≈650 words). Please list the key points you think the author is making in no more than six bullet points.”
- **How I used the output:**
  - Compared the model's six bullets with the three core claims I intended to convey.
  - Noticed that bullet #4 emphasized historical context I meant to downplay; rewrote ¶2 to tighten focus.
  - Confirmed that bullets #1, #2, #5 matched my intended thesis and supporting evidence; no wording from the AI appears in the final paper.
- **Verification:** Re-read the entire paper after revisions to ensure accuracy and coherence.

- **Responsibility:** All ideas and language in the submitted version are my own; the AI served only as a comprehension check.

*My prompt: Students who use PowerPoints have vastly different levels of proficiency with that format. Some make basic errors like trying to fit too much information into one slide, using type that's too small, or using type that doesn't contrast well. Would it be a good use case to upload the draft PowerPoints into generative AI and ask the AI to review the slides for issues like that? And if you think that's a good use case, please draft an AI use statement a student could use if they did this with their PowerPoints.*

### **Is this an appropriate use?**

Yes. Having a model act as a **design-quality reviewer**—flagging crowded slides, illegible fonts, or low-contrast colour schemes—helps students polish visual delivery *without* outsourcing substantive writing or argument. The practice is comparable to running PowerPoint’s own “Accessibility Checker,” but broader. It remains appropriate so long as:

1. **Content ownership stays human.** The AI only critiques layout and readability; no new slide text is accepted from the model.
2. **Student judgment prevails.** The student decides whether each suggestion improves clarity and fits the seminar’s stylistic norms.
3. **Transparency is maintained.** Because the feedback shaped the submitted work, the student discloses it—even though no AI-generated wording appears in the deck.

---

### **Sample AI Use Statement (PowerPoint Design Review)**

#### **AI use statement for presentation slides submitted 3 Nov 2025**

- **Tool & access date:** Microsoft Copilot in PowerPoint (GPT-4o backend), 1 Nov 2025.
- **Purpose:** To identify basic design and accessibility issues (over-full slides, small fonts, low-contrast text/background combinations).
- **Prompt (excerpt):**  
“Review the 12-slide deck I’ve uploaded. For each slide list any problems with information density, font size (<18 pt), or colour contrast below WCAG AA. Do **not** rewrite content; give only diagnostic comments.”
- **Output received:** Copilot produced a slide-by-slide checklist (e.g., “Slide 4: 9 bullet points—consider splitting slide; Slide 7: contrast ratio 3.8:1 between text and background”).
- **How I used the output:**  
– Reduced bullet points on Slides 4 and 6; increased font size on Slide 9; changed background colour on Slide 7 to meet contrast guidelines.

- No AI-generated phrasing or graphics were inserted into the deck.
- **Verification:** Manually ran PowerPoint’s built-in Accessibility Checker and rehearsed the slideshow to ensure legibility from the back of the classroom.
- **Responsibility:** All slide content and design choices in the final version are my own; AI feedback served only as a formatting audit.

*My prompt: I would be fine with the students using generative AI to come up with a paper topic. For example, the student might tell the generative AI, "I want to write about multidistrict litigation but I can't think of a good doctrinal topic that will be suitable for a 25-page paper and that doesn't require competencies that I lack. My competencies are in legal doctrinal research, but not in (for example) empirical analysis or social-science interviews. Please generate a list of possible topics, giving preference to those that are important but can be tackled in the scope of a 25-page paper and focus on a legal doctrinal question." Is this a good use case and, if so, please draft a sample AI use statement for this type of use case.*

### **Is this an appropriate use?**

Yes. Letting a model **brainstorm potential research questions** is analogous to chatting with a librarian or mentor: the student still decides what to pursue and does all subsequent research and writing. It is fine so long as the student:

1. **Owns the final topic choice.** The AI only supplies a menu of ideas; selection and framing are the student’s work.
2. **Checks feasibility.** Before committing, the student confirms that the sources, scope, and skills really match a 25-page doctrinal paper.
3. **Discloses the assistance.** Because the AI output influenced the assignment, the student files an AI use statement—even though no model-generated text appears in the paper.

---

### **Sample AI Use Statement (Topic-Generation Brainstorm)**

#### **AI use statement for final paper topic selection (submitted 18 Sept 2025)**

- **Tool & date used:** OpenAI ChatGPT (model o3), 16 Sept 2025.
- **Purpose:** To brainstorm doctrinally focused, 25-page-scale topics within multidistrict litigation (MDL).
- **Representative prompt:**  
“I want to write about multidistrict litigation but need a legal-doctrinal topic suitable for ~25 pages. My strengths: case law analysis and statutory interpretation. I lack experience in empirical methods. Generate 8–10 potential topics that (a) matter to current MDL practice, (b) turn on legal doctrine, and (c) can be handled without empirical or interview work.”
- **Output received:** A list of nine topic ideas with 2–3-sentence descriptions (e.g., “The constitutionality of ‘rocket-docket’ scheduling orders in MDLs”; “Revisiting the Lexecon

waiver doctrine post-Bristol-Myers”).

• **How I used the output:**

- Screened each idea for novelty and source availability.
- Selected “Lexecon waivers after Bristol-Myers” as my preliminary topic.
- Conducted my own case-law survey to confirm depth of material.
- No wording, phrasing, or citations from the AI list appears in my prospectus or future drafts.

• **Responsibility:** Topic selection was informed by the AI brainstorm, but all framing, research, and writing going forward are exclusively my own.

*My prompt: I need to emphasize to the students that the paper outline needs to reflect THEIR OWN thinking about the paper. I would be very against a student (who was, say, up against the outline submission deadline) simply feeding selected background readings into generative AI and asking it to create an outline. Even if the student then translated the AI's outline into their own language, I would be against this because the whole point of the outlining process is to promote clarity of THEIR OWN thought. Given my view on this, can you think of any acceptable use cases in connection with producing the outline, or is that simply a stage where generative AI use cannot be both appropriate and helpful?*

## Can AI ever be used appropriately at the *outline* stage?

Yes—but only in a **diagnostic/feedback role**, never as a ghost-planner.

Think of the model as an external reviewer who can spot omissions or structural kinks *after* the student has produced an outline generated entirely by their own thinking. Two permissible patterns:

### Permissible “review” use

**Stress-testing logic** – “Here is my outline. List any steps where the inference seems weak or unsupported.”

**Gap spotting** – “Identify doctrinal counterarguments I have not addressed.”

**Organization audit** – “Flag any headings that overlap or could be reordered for better flow.”

**Citation check-list** – “Highlight any claims that would require a case or statute citation.”

### What stays off-limits

Asking the model to *write* or *re-write* the outline, even if the student later re-phrases it.

Feeding background readings and saying “Draft me an outline that synthesises these.”

Copy-pasting AI-suggested headings or bullets into the submitted outline.

Letting the model insert the citations itself (risk of hallucinated sources).

These uses preserve the pedagogical goal—students devise the structure themselves—while giving them a “mirror” to test clarity and completeness.

## Sample AI Use Statement (Outline Diagnostic Review)

### AI use statement for outline submitted 21 Oct 2025

- **Tool & date used:** Anthropic Claude 3.1 (web interface), 18 Oct 2025.
- **Purpose:** To audit my *self-written* outline for logical gaps and redundant sections.
- **Prompt (abridged):**  
“Below is my 1,100-word outline for a 25-page paper on Lexecon waivers post-Bristol-Myers. Please do **not** rewrite any part of it. Instead, point out (1) arguments that lack supporting authority, (2) potential counterarguments not yet addressed, and (3) sections whose order might impede reader flow.”
- **Output received:** A numbered critique highlighting three undeveloped counterarguments and noting that Sections III.A and III.B overlapped.
- **How I used the output:**
  - Added a new sub-heading to engage with sovereign-immunity objections.
  - Merged overlapping sections into a single Part III to streamline progression.
  - No AI-generated wording, headings, or citations appear in the revised outline.
- **Verification:** Cross-checked each flagged gap against case law and manually updated the outline.
- **Responsibility:** The outline’s structure and content remain entirely my own; the AI served solely as a diagnostic consultant.

This disclosure makes clear that AI assisted only in *evaluating* the student’s original outline, not in generating it, thereby upholding your requirement that the outline capture the student’s own thought process.

*My prompt: [I had asked ChatGPT to suggest ways AI might be used in drafting or polishing the final research paper. One of its suggestions was to upload a section of the argument and ask the AI to suggest counterarguments. Here is my prompt following up on that exchange:]*

*I surmise that any LLM that is performing the sort of "evidence check" that you describe here may be sometimes be drawing on humans' previous writings that the LLM has ingested. (Of course, I realize that other times the LLM will be recognizing a pattern that no human has yet recognized, but let's leave that aside for purposes of my question.) Outside the context of AI, I would want the student to cite to that particular piece of prior writing as the source that gave them the idea. I do realize that in our scenario, the "AI use disclosure" can serve a partly similar purpose by showing that the student is acknowledging that this idea is not one they thought of themselves. But the AI use disclosure by itself wouldn't serve the additional goal of citation, which is to give credit where credit is due (in our scenario, to the original human author from whom the AI drew the idea). Is there a guideline that we should add in connection with some uses of generative AI (idea-generating uses, I guess) that would help the student to pick up on whether a prior human author was the originator of the idea, so that the student can credit them in addition to giving the AI use disclosure?*

Below is language you can fold into your syllabus right after the existing “AI use statement” requirement. It gives students a **clear test**—“**Can I trace this idea to a human author?**”—and a **procedure for attribution** that parallels ordinary scholarly practice.

---

### **Proposed Addition: “Source-Tracing for AI-Generated Ideas”**

**If an AI system supplies you with a substantive idea (e.g., a new argument, doctrinal counterpoint, quotation, or historical claim) that you intend to incorporate in your work, you must make a good-faith effort to trace that idea back to a verifiable human source.**

1. **Treat the AI output as a lead, not a citation.**  
*Prompt the model for clues:* “Where in the academic or judicial literature has this argument appeared? Please list specific cases, articles, or books.”
2. **Independently confirm any sources named.**
  - Look up the case, article, or book yourself.
  - Read enough of the source to be sure it actually contains the idea.
3. **Cite the *human* source in the ordinary way.**
  - Example: *See* Smith, **The Post-Lexecon Landscape**, 54 U. Chi. L. Rev. 321, 334–35 (2024).
4. **If, after reasonable search, no prior human source emerges**, state in your AI use statement that:  
“The model suggested Argument X; I searched the legal literature (databases A, B) and was unable to locate a prior publication. The argument therefore appears to be original to me, though inspired by the AI prompt.”
5. **Never cite the AI model itself as the intellectual originator** of a legal or scholarly claim. Models provide *computations over text*; citation credit belongs to the human author(s) on whose text the model drew—or to you if no prior source exists.

---

### **Rationale to Share with Students**

- **Credit where credit is due.** Citing the underlying human author satisfies the scholarly norm that readers should be able to trace an idea to its first articulation.
- **Academic integrity.** The AI use statement discloses *process*; formal citations document *provenance*. Both are required for full transparency.

- **Error control.** Verifying the source guards against hallucinated citations and misattributed ideas.
- 

### Quick Reference Checklist for Students

1. **Did the model give me a substantive idea I hadn't already formed?**
2. **Can the model (or a database search) point me to a human author?**
3. **Have I read and confirmed that source?**
4. **Have I cited that source in the paper/footnotes?**
5. **Have I described the AI's role in my AI use statement?**

Including this guideline will help students distinguish between permissible brainstorming assistance and the scholarly obligation to acknowledge the real intellectual lineage of the ideas they deploy.

*My prompt: great. so now let's draft a sample AI use statement (plus citation to human author) for use in this scenario: the student asked the AI to point out plausible doctrinal counter-arguments on one of the student's arguments, and then when the AI offered such a counter-argument, the student followed up with further prompts to try to trace the counter-argument to a human source, and did find the source.*

### Sample AI Use Statement + Scholarly Citation

*(Model a student might append to the end of the paper, immediately after the standard bibliography)*

#### **AI use statement for final research paper (submitted 15 Dec 2025)**

- **Tool & date:** OpenAI ChatGPT (model o3), accessed 8 Dec 2025.
- **Purpose:** To stress-test Section III of my draft, which argues that post-*Bristol-Myers* personal-jurisdiction limits should not constrain transferee courts from approving *Lexecon* waivers.
- **Prompts (abridged):**
  1. "Here is Section III of my draft ( $\approx$ 1,400 words). List any plausible doctrinal counter-arguments a court might raise."
  2. *Follow-up:* "For Counter-Argument #2 you suggested— that *Lexecon* waivers could violate defendants' Seventh Amendment rights—where (if anywhere) has this claim appeared in published scholarship or case law? Please name specific sources."
- **Output received:**
  - Five counter-arguments identified.
  - Cited **Professor Dana M. Yale**, "**Seventh Amendment Constraints on Lexecon Waivers**,"

72 *Fordham L. Rev.* 945 (2004) as prior articulation of Counter-Argument #2.

• **How I used the output:**

1. Read Yale’s article in full to verify the argument’s scope.
2. Added a new subsection, § III B 2, responding to Yale’s Seventh Amendment critique.
3. All language in the paper is my own; no AI-generated wording appears verbatim or in paraphrase.

• **Verification:** Confirmed Yale’s page numbers and doctrinal analysis; double-checked Bluebook citation format.

• **Responsibility:** AI served solely as a diagnostic aid to locate existing scholarship. The counter-argument’s substantive discussion and all responses are my original work.

---

### Example in-text treatment of the human source

Some scholars contend that allowing parties to waive § 1407’s remand right could impermissibly burden defendants’ Seventh Amendment jury guarantee.<sup>1</sup> I argue below that this concern misreads both *Beacon Theatres* and modern MDL practice.

---

### Corresponding footnote

<sup>1</sup> Dana M. Yale, *Seventh Amendment Constraints on Lexecon Waivers*, 72 *Fordham L. Rev.* 945, 960–62 (2004).

---

This model disclosure:

- **Separates the AI’s role** (idea lead) from the **human author’s role** (intellectual originator of the counter-argument).
- Shows a **trace-and-verify step** before citing.
- Credits the source in standard scholarly form while the AI use statement documents process transparency.