



ARTICLE

Incorporating AI into the Contract Drafting Process: A Classroom Exercise

Kristen V. Hardy

Senior Legal Skills Professor

University of Florida Levin College of Law

Published: December 2025

I recently read the *Three Blind Drafts: An AI-Generated Classroom Exercise*,¹ by Margie Alsbrook and Ashley Chase. I was trying to create an AI-related assignment for my spring contract drafting class. The article inspired me to attempt a conceptually similar assignment in an AI contract drafting exercise. This article will summarize Alsbrook's and Chase's original idea and describe how I adapted it for the contract drafting classroom; the student outcome; and my thoughts about what I would do differently the next time I attempt the assignment.

I. The Original Assignment

In their article, Professors Alsbrook and Chase detailed their creation of a classroom exercise that involved getting their objective legal writing students to use generative AI tools that are currently available in the marketplace. The stated outcome of the exercise was that "students quickly grasp

¹ Margie Alsbrook & Ashley Chase, *Three Blind Drafts: An AI-Generated Classroom Exercise*, 37:3 *The Second Draft* (Spring 2025), <https://www.lwionline.org/article/three-blind-drafts-ai-generated-classroom-exercise> (last visited May 7, 2025).

pitfalls of the tools, while they also start to understand that different AI products suit different purposes.²

The professors created a sample fact scenario and prompt for a research memo and gave the assignment to their students. After the student memos were completed, the professors fed an identical prompt based on that scenario to each of the three generative AI sources: ChatGPT, Lexis+ AI, and Claude.³ Each of the AI sources created a sample memo, which the professors then shared with the students.⁴ The professors did not tell the students that the samples were the products of artificial intelligence.⁵ The students offered critiques of the three writing samples, guided, in large part, by reflection questions posed by the professors.⁶ Then the professors revealed that ChatGPT, Lexis+ AI, and Claude had drafted the samples. This led to a lively discussion about the role of AI in legal research and writing.⁷ And the students concluded that, “while AI can assist in generating drafts and organizing thoughts, it cannot replace the need for independent verification of sources and a deep understanding of the law.”⁸

2. The Contract Drafting Course Revamp: The Prelude

I want to be very deliberate about how I instruct students on the use of AI in my contracts drafting course. Some of my colleagues embrace AI wholeheartedly and allow students to use it throughout the course. Other colleagues forbid students from discussing AI in class, let alone use it on assignments. My approach organically falls somewhere in the middle. I confess to students that I am still learning about generative AI and its uses, particularly in the contract drafting space. However, I think there are wonderful applications for AI, particularly as a first draft or checklist generator. I caution students about the limitations of generative AI, though, and emphasize the importance of incorporating AI into the drafting process *after* first solidifying an understanding of the mainstream drafting

² *Id.*

³ *Id.* The professors shared that in “later experimentations” they learned that the AI sources generated better prompts if they slightly altered the prompt “to account for the strengths and weaknesses of each AI product.”

⁴ *Id.*

⁵ *Id.* at 2-3.

⁶ *Id.* at 3.

⁷ *Id.*

⁸ *Id.*

concepts.⁹ So, this past spring semester I set out to craft an AI exercise that embraced the capabilities of the technology, but that also displayed the limitations inherent in its inability to draft according to the conventions and concepts we learned in the course and in our textbook.

3. The In-Class Drafting Exercise

4.

For this assignment, I had students revisit a prompt and set of deal terms that they had worked with earlier in the semester. I felt that one of the key features of the original assignment was that students were familiar with the fact scenario, so I wanted to make sure I replicated this feature.

Prior to the class session, I ran the prompt (“I’d like you to draft a contract based on the following deal terms and facts”) and deal terms through three large language models (LLMs): ChatGPT 4.0, Lexis+ AI Protege, and Claude 3 Opus. Each LLM produced a sample agreement based on the instructions provided.

Each LLM-generated sample was wildly different from the others. I didn’t conduct this exercise blind, as was the case with the original assignment. Students were aware that we were going to explore the use of AI in contract drafting. It was on the syllabus. At the start of our class time, I explicitly told students that the samples were AI-generated and told them which LLM was used to draft each sample.

I instructed students to take the final draft that we had written as a class (according to the conventions in our textbook) and use it as a standard against which to “grade” the AI-generated sample agreements. To guide their analysis, I put seven questions on the board for students to use to help assess the merits of each sample.

3.1 Logistics

This assignment contained a lot of moving parts; so for functional ease, I uploaded all assignment components into a Canvas module. To complete the exercise, students had the following documents before them:

⁹ *Id.* at 2 (One key point that I took away from the article written by Professors Alsbrook and Chase was that “[t]he exercise works best when students are already familiar with the legal questions and issues that apply to the sample factual scenario.”)

1. The In-Class [Collaborative] Assignment instructions (including deal terms);
2. The In-Class [Collaborative] Assignment final draft (working copy);
3. Sample agreement generated by ChatGPT 4.0;
4. Sample agreement generated by Lexis+ AI Protege; and
5. Sample agreement generated by Claude 3 Opus.

I also created drop boxes for the groups of students in each class to submit their assignments.

Week 11: April 1 and 3	✓	+	⋮
Week 11 Module	✓	⋮	⋮
+Merchandise Licensing Agreement - Working Copy.docx	↻ ✓	⋮	⋮
(INSTRUCTIONS FOR AI) S25 Licensing Contract (deal terms).docx	↻ ✓	⋮	⋮
Chat GPT 4.docx	↻ ✓	⋮	⋮
Claude 3 Opus.docx	↻ ✓	⋮	⋮
Lexis+ AI Protege.docx	↻ ✓	⋮	⋮
AI In-Class Group Assignment (April 1, 2025) - One Submission Per Group (C674) 0 pts	✓	⋮	⋮
AI In-Class Group Assignment (April 1, 2025) - One Submission Per Group (C677) 0 pts	✓	⋮	⋮

3.2 Questions for Student Evaluation

Professors Alsbrook and Chase advised that “[a]n important part of this AI classroom exercise is the selection of evaluation and reflection questions.”¹⁰ With that guidance, I tailored the reflection questions in a way that would move the conversation about the usefulness of AI in the drafting process forward, and answer, definitively for students, whether the LLM reviewed held any value. I used the following questions in the classroom exercise:

1. In thinking about the concepts and techniques we have studied in this class, list at least three things you think the sample does well

¹⁰ *Id.* at 3.

and why. Then list at least three things you think the sample does not do well and why.

2. Did the sample capture all relevant
 - a. covenants, and
 - b. representations and warranties?
3. Did the sample organize the provisions in a way that is consistent with the parts of the contract list we have studied in class?
4. Are there any key provisions that were left out of the sample?
5. Are there any provisions in the sample that you would strike altogether?
6. How would you improve the instructions to generate a better sample?
7. Would you recommend this generative AI tool for drafting contracts? Why or why not?

3.3 The Evaluation Process

After posting the reflection questions on the board, I handed out hard copies of the three AI-generated samples and published them on our Canvas page. I allowed my students to break off into groups and gave each group¹¹—twenty minutes to work through the reflection questions for each sample.

I taught two sections of this class, back-to-back, and gave students slightly different instructions at this point in the assignment. In my first class, it was clear that twenty minutes wasn't enough time to get through all of the reflection questions for each sample (I'll revisit this issue later), and most

¹¹ Because I incorporate a lot of drafting exercises into our class time, I customarily have students organize themselves into groups of no more than four students. I usually have them do this in the second class of the semester. Although we don't do group work every day, we do it often enough that when that time comes, students already know who they will be working with on the in-class assignment.

groups had not gotten to the Lexis+ AI Protege sample, which was the third sample on the list. So, I quickly adjusted when I taught the second section later in the day. With this second class, I assigned each group a different sample to start with so that I could guarantee that at least one group had evaluated each sample, which allowed us to have a more robust discussion on the comparative virtues of each sample.

3.4 Student Feedback

Students shared their feedback in class when we reconvened after twenty minutes of group work, and I asked them to upload their comments to a Canvas drop box so I could retain them. The students had a general consensus about the value of each sample when compared with the other samples and against the assignment's model answer. One particular group's written reaction was representative of the whole class, as the other groups shared some variation of the following feedback:

LLM #1: ChatGPT:

Question 1: In thinking about the concepts and techniques we have studied in this class, list at least three things you think the sample does well and why. Then list at least three things you think the sample does not do well and why.

Three things the sample did well:

1. It designated parties correctly (excluding the bolding which maybe was lost in the download process).
2. It included the correct sections at the beginning (definitions, preamble, etc.).
3. It tried to define a decent amount of terms.

Three things the sample did not do well:

1. It labeled provisions unnecessarily like the preamble.
2. The sample used legalese.
3. It [didn't number the] articles.

Question 2: Did the sample capture all relevant covenants, and representations and warranties?

The sample does not capture all relevant covenants and even worded them wrong (making the covenants the exact opposite as in the instructions).

The sample does not capture all relevant representations and warranties. It did not even phrase them properly.

Question 3: Did the sample organize the provisions in a way that is consistent with the parts of the contract list we have studied in class?

The sample generally organized the provisions like we studied in class.

Question 4: Are there any key provisions that were left out of the sample?

There are key provisions left out in the sample. There are no endgame provisions.

Question 5: Are there any provisions in the sample that you would strike altogether?

There were no provisions that we would strike through altogether, but we would change how they were worded.

Question 6: How would you improve the prompt to generate a better sample?

We would improve the prompt to get a better sample by adding more information about structure or even uploading a sample that conforms to what we like.

Question 7: Would you recommend this generative AI tool for drafting contracts? Why or why not?

We would not recommend this generative AI tool for drafting contracts because it has no formatting ability and the contract reads almost like a stream of consciousness in that it has no structure.

LLM #2: Claude 3 Opus:

Question 1: In thinking about the concepts and techniques we have studied in this class, list at least three things you think the sample does well and why. Then list at least three things you think the sample does not do well and why.

Three things the sample did well:

1. It designates parties.

2. It includes entity information.
3. The termination provision is pretty solid.

Three things the sample did not do well:

1. It uses legalese.
2. There are no definitions.
3. There is no numbering.

Question 2: Did the sample capture all relevant covenants, and representations and warranties?

The sample does not capture all relevant covenants, as it excludes the territory information. The sample does not capture all relevant representations and warranties, as the ones it does provide are cursory and non-specific.

Question 3: Did the sample organize the provisions in a way that is consistent with the parts of the contract list we have studied in class?

The sample did not organize the provisions like we studied in class.

Question 4: Are there any key provisions that were left out of the sample?

There are key provisions left out in the sample such as definitions and payment provision.

Question 5: Are there any provisions in the sample that you would strike altogether?

There are provisions that we would strike through altogether, such as the miscellaneous provision. We would instead remove this and add the provisions under it into separate provisions.

Question 6: How would you improve the prompt to generate a better sample?

We would improve the prompt to get a better sample by uploading a sample.

Question 7: Would you recommend this generative AI tool for drafting contracts? Why or why not?

We would not recommend this generative AI tool for drafting contracts because it excludes major portions of the contract, organizes poorly, and is very hard to read.

LLM #3: Lexis+ AI Protégé:

Question 1: In thinking about the concepts and techniques we have studied in this class, list at least three things you think the sample does well and why. Then list at least three things you think the sample does not do well and why.

Three things the sample did well:

1. It numbers the provisions.
2. It designates the parties.
3. Some internal organization of the articles.

Three things the sample did not do well:

1. It uses legalese.
2. The definitions are not in alphabetical order.
3. Recitals are too long.

Question 2: Did the sample capture all relevant covenants, and representations and warranties?

The sample does capture all relevant covenants, even including the territory expansion provision. The sample does not capture all relevant representations and warranties, excluding licensee representations.

Question 3: Did the sample organize the provisions in a way that is consistent with the parts of the contract list we have studied in class?

The sample does generally organize the provisions like we studied in class, just the recitals are a bit odd in formatting.

Question 4: Are there any key provisions that were left out of the sample?

There are key provisions left out in the sample, such as the licensee's representations.

Question 5: Are there any provisions in the sample that you would strike altogether?

There are not provisions that we would strike through altogether, however, we would reorganize some like moving up the term provision towards the top of the contract.

Question 6: How would you improve the prompt to generate a better sample?

We would improve the prompt to get a better sample by uploading an example, give it more guidance on endgame provisions, and tell it to remove legalese.

Question 7: Would you recommend this generative AI tool for drafting contracts? Why or why not?

We would sort of recommend this generative AI tool for drafting contracts because it is organized more like a real contract, but it still has enough issues that drafting it yourself is best. However, of the tools we explored, this one is the best. It could provide a basic framework to use, but I would not rely on it.

In summary, students identified similar issues with all three samples: omission of important information from the fact pattern like the payment provisions, certain representations and warranties, and the termination threshold. Most determined that the sample produced by Lexis+ AI Protege most closely aligned - in format and in content - with the way they had been taught to draft in the course. Some students offered that one way they would try to get a better sample would be to go further than to just upload the instructions, as I had done, and actually upload a sample contract from our class to give the LLM guidance on expected format and content.

3.5 Suggested Edits for the Next Time I Use the Assignment

As I did with my second section of the day, I plan to revise and streamline the reflection questions. I also plan to give students a longer amount of time to work on the exercise. I didn't know how well the exercise would execute, so I scheduled it on a day that the students had to complete course evaluations, which ate up the first fifteen minutes of class. Because

I know that students want more time to work with the AI-generated samples, I will plan to have my students complete this exercise on a day in which they have the full class period to complete.

5. Conclusions

The decision to tell students in advance that the samples would be generated by artificial intelligence was an important one. By going into the exercise with this awareness, students were better able to focus their critique on the value of AI as a tool. Viewing AI as a tool, and not a substitute for lawyering skills, is key.

Therefore, this assignment holds value for three main reasons. First, it pushes back on the myth that AI will replace lawyers. When it comes to contract drafting (more so than memo and brief writing), students often assume that they can make a decent effort by stringing together boilerplate provisions. This assignment reminds us of the necessity of human contribution to the drafting process. AI can be a useful tool for organizing thoughts and for generating a checklist of important provisions to include, but it is not a substitute for the careful drafting of each provision that is required of every legally binding agreement. In the contract drafting space, competent client representation goes beyond the simple inclusion of a list of provisions, but requires the attorney to carefully craft the contract to be accurate, clear, concise, and to mitigate their client's risk of exposure to future litigation.

Second, the mediocre quality of the AI-generated samples suggests that a young attorney's time might be better spent applying the drafting conventions they are familiar with and drafting the agreement on their own. The purpose of using AI is to save time.¹² So, it is counterproductive to spend so much time educating a large language model enough to generate an output that will need even more editing and revision.

Lastly, this assignment underscores the fact that an agreement that takes all parties' nuanced interests into account is a product of a collaborative process. Attorneys who know the law and understand their clients' needs are best positioned to collaborate on a mutually beneficial outcome. Even

¹² Huffington, Ariana, "AI will save us time. The real question is what we'll do with it," <https://fortune.com/2025/10/24/ai-will-save-us-time-the-real-question-is-what-well-do-with-it-arianna-huffington/> (last accessed on December 16, 2025).

though attorneys may be on opposite sides of the contract drafting table, they share the objective of solidifying a deal that is grounded in coordination, flexibility, patience, communication, and compromise. Everyone wants the deal done, and the process of getting to “done” requires real collaboration and can’t be accomplished through simple cut-and-paste or the use of imprecise standard provisions. Regardless of whether and to what extent AI is incorporated into the contract drafting process, by placing AI in its proper place, students are better equipped to navigate the digital literacy expectations of their future employers.