

5-10-2024

## Artificial Intelligence: Legal Reasoning, Legal Research and Legal Writing

S. Sean Tu

Amy Cyphert

Samuel J. Perl

Follow this and additional works at: <https://scholarship.law.umn.edu/mjlst>



Part of the [Legal Writing and Research Commons](#), and the [Science and Technology Law Commons](#)

---

### Recommended Citation

S. Sean Tu, Amy Cyphert & Samuel J. Perl, *Artificial Intelligence: Legal Reasoning, Legal Research and Legal Writing*, 25 MINN. J.L. SCI. & TECH. 105 (2024).

Available at: <https://scholarship.law.umn.edu/mjlst/vol25/iss2/11>

# Artificial Intelligence: Legal Reasoning, Legal Research and Legal Writing

S. Sean Tu \*, Amy Cyphert\*\*, & Samuel J. Perl\*\*\*

## INTRODUCTION

The key elements of a lawyer's work can be broken into three component parts: legal search, creating legal arguments, and legal writing. Artificial Intelligence (AI) can play an important role in each of these steps.

Legal search, the first key step in legal research, helps the lawyer understand the current black letter law and understand how the facts of her case may differ from current legal precedent. Legal search involves finding those cases that are most factually similar to the instant case. If a prior case is factually identical to the instant case, and from the same jurisdiction, little further analysis will be required. By retrieving the similar cases based on the facts of the case and other easily identified features (such as jurisdiction), AI methods using *supervised* learning are useful methods for performing this task.

However, legal analysis is not based solely on similar factual situations. Crafting legal arguments is possibly the most important step for a lawyer and requires the understanding of the public policy behind the law and the foundational legal theories that created the law. If the facts are not identical to a previous case, then understanding what facts are most important to distinguish a previous ruling, or to analogize to a previous ruling that is beneficial to the client's position, can be the difference between winning or losing. This step requires understanding not only the key factual differences but also the underlying legal theory that is being used, as well as the role of public policy behind the law or previous precedential cases. Creating a legal argument requires analogical reasoning; AI's *unsupervised* or even *self-supervised* learning may increasingly play an important role in helping lawyers create winning legal arguments.

Finally, crafting a legal memo or brief is the culmination of the lawyer's legal research and legal arguments.<sup>1</sup> Crafting a well-organized, well-written brief that highlights the key facts, issues, legal rules, and public policy arguments is crucial for producing effective descriptive and advocacy work. Generative AI models are increasingly using a technique called *self-supervised learning*. While AI may play an important role in creating first drafts, it has the potential to impede the lawyer's creative process and could be detrimental to the profession if relied upon too heavily.

This article is organized in three parts. Part I describes Large Language Models (LLMs), both the underlying technology and how they are currently being used. Part II describes legal search and how AI might be used to increase efficiency and find better cases based on similar facts. Part III describes the creation of legal arguments and how AI's unsupervised learning may increasingly play a role in helping a lawyer create legal arguments based on the legal theory and public policy underpinning the law and previous precedential opinions or statutes. Part IV outlines the legal writing process and describes how generative AI can be a significant asset in assisting lawyers with crafting a conclusive legal memo or advocacy brief.

## I. ARTIFICIAL INTELLIGENCE AND LARGE LANGUAGE MODELS

### A. EVOLUTION OF AI TOOLS

Researchers and inventors have been attempting to build language generating tools for decades, with many small successes along the way. For example, in 1966, MIT researcher Joseph Weizenbaum created a computer program which made

---

© 2024 S. Sean Tu JD, PhD, Amy Cyphert, and Samuel J. Perl

\* West Virginia University College of Law, Morgantown, WV.

\*\* West Virginia University College of Law, Morgantown, WV. This article was supported by a Hodges Fund faculty research grant.

\*\*\* Carnegie Mellon University, Pittsburgh, PA.

1. Of course, each of these three component parts are interrelated and a litigator will often move back and forth between them. For example, as is discussed further on, it is not unusual for an attorney to realize in the process of drafting their brief that they have a logical problem with their legal reasoning, and thus need to go back to do further research and refine their arguments.

possible certain communications between people and machines.<sup>2</sup> His goal was to allow humans to have conversations with machines using natural language. The machine would recognize the parts of language that the users inputted and could provide reasonable and appropriate responses, albeit in a limited and rudimentary fashion.<sup>3</sup> There have since been many other efforts in the field of Natural Language Processing (NLP) to develop tools with better capabilities for categorizing and outputting human language.

The latest generation of AI tools that generate written text are commonly referred to as Large Language Models (LLMs). These models use techniques from the field of deep learning to create AI tools capable of generating written text that mimics human speech. OpenAI is recognized as a leader in this field, having created the well-known and popular LLM ChatGPT. To create the Generative Pre-Trained Transformer 3 (GPT-3) Language Model, the precursor to ChatGPT, OpenAI collected or scraped large amounts of text from the internet and used deep learning techniques to train a new language model that could predict the next word in a sentence. GPT-3 amazed its initial users when it was released in 2020, but it had some limitations when users wanted it to be more specific to their needs.<sup>4</sup> AI companies, including OpenAI, took notice of these limitations and designed techniques to improve the ability for users to direct the LLM to generate the text that was more specific to their needs. Companies also added the ability for users to provide instructions directly to the language model to improve user satisfaction with its output using a process that OpenAI calls post-training alignment.<sup>5</sup>

---

2. Joseph Weizenbaum, *ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine*, 9 COMM'NS ACM 36, 36 (1966).

3. *Id.* at 36.

4. Rob Toews, *GPT-3 Is Amazing—And Overhyped*, FORBES (July 19, 2020, 6:56 PM), <https://www.forbes.com/sites/robtoews/2020/07/19/gpt-3-is-amazingand-overhyped/> (noting that GPT-3 answered the question “How many eyes does my foot have?” with “Your foot has two eyes.”).

5. *Aligning Language Models to Follow Instructions*, OPENAI (Jan. 27, 2022), <https://openai.com/research/instruction-following> (“GPT-3 models aren’t trained to follow user instructions. Our InstructGPT models [] generate much more helpful outputs in response to user instructions.”).

Multiple models and AI training techniques were used to achieve these features.<sup>6</sup> Users have responded very favorably to the newest generation of LLMs. Analysts at the Swiss bank UBS estimated that ChatGPT reached over 100 million monthly active users just two months after it launched.<sup>7</sup> OpenAI and other AI firms (including Meta, Anthropic, and Google) have launched even more improved models. For example, OpenAI currently makes GPT-4 available to users who pay a monthly subscription fee. GPT-4 was trained using even more parameters than GPT-3 or ChatGPT.<sup>8</sup> GPT-4 also expanded how users can give the model instructions, allowing users to submit different types of file inputs including images, sounds, music, and more.

## B. THE FUTURE OF AI TOOLS FOR LAW?

Today marks a new age of AI tools for lawyers. Not only are AI-powered information retrieval and analysis techniques being used for legal tasks, including supervised and unsupervised learning, but generative AI powered by pre-trained language models holds new promise.

Pre-Trained Language Models now contain representations of different human concepts learned from their large training sets.<sup>9</sup> They are also amazingly good at predicting reasonable and grammatically correct words in certain contexts.<sup>10</sup> They are not without their limitations, including hallucinations (AI-generated false or misleading information), bias, copyright concerns, and client confidentiality issues, but the possible benefits of generating reasonably usable legal documents at lowered costs are exciting. Advances in AI alignment, which

---

6. *Introducing ChatGPT*, OPENAI (Nov. 30, 2022), <https://openai.com/blog/chatgpt>.

7. Krystal Hu, *ChatGPT Sets Record for Fastest-Growing User Base – Analyst Note*, Reuters (Feb. 2, 2023, 9:33 AM), <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>.

8. OPENAI ET AL., GPT-4 TECHNICAL REPORT 1 (arXiv 2024) (“GPT-4 is a Transformer-based model pre-trained to predict the next token in a document. The post-training alignment process results in improved performance on measures of factuality and adherence to desired behavior.”). *See also GPT-4*, OPENAI (Mar. 14, 2023), <https://openai.com/research/gpt-4>.

9. Christopher D. Manning, *Human Language Understanding & Reasoning*, 151 DAEDALUS 127 (2022), <https://www.amacad.org/publication/human-language-understanding-reasoning>.

10. *Id.* at 134.

aims to guide AI systems towards a user's intended goals, preferences, ethical standards, and safety has also improved in recent years.<sup>11</sup> AI alignment will improve even faster as more generative AI models are released and refined with real users' data as companies learn directly what users want (and do not want).

Pre-Trained Language Models allow companies that want to build AI tools for lawyers to start with a model that already understands human language, has a vast amount of knowledge learned from text, and is already partially aligned to human user feedback when generating its output. Organizations are already working with lawyers to align models even more specifically to their needs and contexts. Thus, models do not need to be built from scratch anymore. Nevertheless, there may be a need for more effective application of the model.

### C. BASICS OF SUPERVISED LEARNING, UNSUPERVISED LEARNING, AND PRE-TRAINED LANGUAGE MODELS (SELF-SUPERVISED LEARNING)

#### 1. Supervised Learning

The goal of supervised learning is often to predict outcomes. This is achieved through guiding the algorithm with labeled input and correlating those known inputs to known outputs. In this process, the algorithm is trained using data with pre-existing known outcomes, where both the input and output are labeled. Consequently, the training data serves as a teacher for the algorithm.

The instructional process in supervised learning involves several steps. First, the training data incorporates as many examples as possible of the subject being taught. Second, there is a selection of the model's architecture (transformer architecture, convolutional neural networks, or recurrent neural networks are three popular choices). Third, a penalty, also known as a loss function, is created, which the algorithm uses to assess the degree of damage for an erroneous decision, thus facilitating learning for greater accuracy in subsequent attempts. Based on these results, the model updates its

---

11. For example, reinforcement learning techniques are used to align models after training to human feedback. See Nathan Lambert et al., *Illustrating Reinforcement Learning from Human Feedback (RLHF)*, HUGGING FACE (Dec. 9, 2022), <https://huggingface.co/blog/rlhf>.

algorithm to reduce the loss function, resulting in increased accuracy of the prediction. To do this, the model employs technologies like Stochastic Gradient Descent (SGD) and backpropagation. The process is termed supervised learning because the training data contains labels that are both used to make predictions and to indicate if the predictions the model made were accurate. Thus, the model has supervision in the form of correct and incorrect answers.

## 2. Unsupervised Learning

Where the objective of supervised learning is usually to predict outcomes, the objective of unsupervised learning is usually to uncover inherent patterns and extract valuable insights. This approach involves providing the algorithm with a set of unlabeled inputs. Successful unsupervised learning algorithms will autonomously explore patterns and trends in the unlabeled inputs to identify latent relationships. In unsupervised learning, the algorithm is trained on unlabeled data without specific guidance, needing to discern the output independently. It essentially groups items with similar characteristics, identifying patterns and trends within the dataset. This form of learning is applicable to addressing association and clustering challenges.

Unsupervised learning techniques can uncover patterns that were previously hidden to humans. For example, unsupervised learning techniques have been used to identify groups of similar customers (customer segmentation),<sup>12</sup> rank webpages (PageRank),<sup>13</sup> partition medical patients with similar disease characteristics,<sup>14</sup> and even measure text similarity between different court case reports.<sup>15</sup>

---

12. Boyu Shen, *E-Commerce Segmentation via Unsupervised Machine Learning*, in CONF-CDS 2021: THE 2<sup>ND</sup> INTERNATIONAL CONFERENCE ON COMPUTING AND DATA SCIENCE 45 (2021).

13. See TREVOR HASTIE ET AL., *THE ELEMENTS OF STATISTICAL LEARNING* 576 (2d ed. 2009) (describing Google's PageRank algorithm).

14. Hany Alashwal et al., *The Application of Unsupervised Clustering Methods to Alzheimer's Disease*, FRONTIERS COMPUTATIONAL NEUROSCIENCE, May 2019, at 1.

15. Arpan Mandal et al., *Unsupervised Approaches for Measuring Textual Similarity Between Legal Court Case Reports*, 29 A.I. & L. 417 (2021).

### 3. Pre-Trained Language Models (Self-Supervised Learning)

Pre-Trained Language Models (including LLMs) use a combination of supervised and unsupervised training techniques frequently called self-supervised learning. The goal of self-supervised learning is to create models capable of learning representations from an unlabeled dataset (aka unsupervised). Self-supervised learning requires large amounts of high-quality (typically human-created) types of data. The technique has been used with success in different fields including computer vision and natural language processing.<sup>16</sup>

Most generative AI models use self-supervised learning techniques combined with a post-training alignment process, such as Reinforcement Learning with Human Feedback,<sup>17</sup> Constitutional AI,<sup>18</sup> Direct Policy Optimization,<sup>19</sup> and others.<sup>20</sup> Researchers are also experimenting with techniques to solve alignment and safety related concerns including reducing hallucinations for certain tasks,<sup>21</sup> adding the ability to cite sources,<sup>22</sup> performing more concise text summaries,<sup>23</sup> and

---

16. Alec Radford et al., *Learning Transferable Visual Models from Natural Language Supervision*, in INTERNATIONAL CONFERENCE ON MACHINE LEARNING 1 (2021).

17. Long Ouyang et al., *Training Language Models to Follow Instructions with Human Feedback*, in 36 NEURAL INFORMATION PROCESSING SYSTEMS 1 (2022).

18. YUNTAO BAI ET AL., CONSTITUTIONAL AI: HARMLESSNESS FROM AI FEEDBACK 1 (arXiv 2022).

19. Rafael Rafailov et al., *Direct Preference Optimization: Your Language Model is Secretly a Reward Model*, in 37 NEURAL INFORMATION PROCESSING SYSTEMS 1 (2023).

20. Sebastian Raschka, *LLM Training: RLHF and Its Alternatives*, AHEAD OF AI (Sept. 10, 2023), <https://magazine.sebastianraschka.com/p/llm-training-rlhf-and-its-alternatives> (summarizing other technical AI alignment techniques). See also Tom B. Brown et al., *Language Models are Few-Shot Learners*, in 33 NEURAL INFORMATION PROCESSING SYSTEMS 1 (Hugo Larochelle et al. eds. 2020) (explaining of Deep Learning Techniques, Transformer Architectures, Large (Billions of Parameters), and Text Tokenization techniques).

21. HUNTER LIGHTMAN ET AL., LET'S VERIFY STEP BY STEP 1 (arXiv 2023).

22. *Gemini Apps FAQ*, GOOGLE, <https://bard.google.com/faq> (last visited Apr. 12, 2024) (“If Gemini does directly quote at length from a webpage, it cites that page. For answers with URLs or image thumbnails, Gemini enables users to easily see and, in some cases, click to navigate directly to the source for each.”).

23. *Release Updates 2024.03.04*, GOOGLE, <https://bard.google.com/updates> (last visited Apr. 12, 2024).



allowing users to generate longer responses.<sup>24</sup> Generative AI could prove a promising technology in law, provided its responses are carefully reviewed, edited, and any citations properly sourced by its users.

#### D. TECHNOLOGICAL ISSUES ASSOCIATED WITH APPLICATION OF THE AI TOOL ON LEGAL RESEARCH AND WRITING

Researchers have identified many issues with the use of AI tools like LLMs, including the propensity of these tools to hallucinate or to make up data, including court cases. This is obviously a serious problem for legal search, though it is important to note that the companies developing these LLMs explicitly state that they are prone to hallucinations and thus should not be used for tasks like legal research.

AI researchers use the term alignment to describe the goal of an AI system doing what the user wants it to do or being aligned with the user's intentions. The term is also used in a broader sense to describe AI that is aligned with the best interests of society. Currently, there is an alignment gap between the performance of certain AI systems and the needs of lawyers. Beyond AI hallucinations, generative AI models may be overly broad, make mistakes in citation, not understand the distinction between precedential and binding authority, and not be deployed in a way that allows lawyers to safeguard client confidentiality.

## II. LEGAL SEARCH

Legal search is the first key step in legal research. A lawyer must find judicial precedents, statutes, case law, regulations, and secondary material that may be relevant to the case at hand. Typically, legal search involves using legal tools and databases (e.g., Lexis Nexis or Westlaw) to efficiently locate pertinent legal documents and precedents. Each field may have specialized tools to help with legal searching. For example, in the patent field, Lex Machina can be used to search patent litigation and Patent Advisor can be used to search patent file histories.

Legal analysis cannot be divorced from legal reasoning. It is "impossible to do legal research without analyzing, synthesizing,

---

24. *Controlling the Length of OpenAI Model Responses*, OPENAI, <https://help.openai.com/en/articles/5072518-controlling-the-length-of-completions> (last visited Apr. 12, 2024).

and applying the information found, both to the original issue and to the research plan developed to address the issue.”<sup>25</sup> At its heart, legal research involves finding cases that are similar to the instant factual situation so that the lawyer can make precedential or analogic arguments (and the related work of distinguishing any cases that might be binding on the court but do not support the client’s position). But more than that, legal research is an interactive process of problem-solving, requiring legal reasoning and analysis.

To help analyze the interaction between legal research and legal reasoning, legal searching is broken into (1) the same field or (2) the same problem. The same field search is a typical keyword search that is based on supervised learning and might return results typically based on factual labeling or pre-cultivated organization of data. This type of search would be most relevant for precedential reasoning. The same problem type of search would be based on unsupervised learning and might return results based on legal principles. This type of search would be most relevant for analogical reasoning.

#### A. OVERARCHING HYPOTHETICAL

A hypothetical scenario will help explain how AI might be used in each of the legal processes (i.e., legal search, legal reasoning, and legal writing). This hypothetical scenario will use a simplified version of the search and seizure rule and serially modify the factual scenarios to illustrate where and how AI might be used as a legal tool, thereby shedding light on its potential applications, benefits, and limitations.

For the search and seizure hypothetical, two simple rules are assumed:<sup>26</sup> (1) a police officer must have a search warrant to search a home, and (2) an officer does not need a search warrant to search an automobile. As an initial matter, for those cases that clearly involve a defendant’s home or automobile, the only legal analysis that may be necessary is to find a case that has identical facts. Also assume that the public policy rationale for these rules is based on a person’s right to privacy. Specifically, a

---

25. Sarah E. Valentine, *Legal Research as a Fundamental Skill: A Lifeboat for Students and Law Schools*, 39 U. BALT. L. REV. 175, 210 (2010).

26. Case law in this field is more complicated than these two rules. They are offered here only for purposes of setting up the hypothetical and not as objectively true statements of law that are universally applicable.

defendant has a clear right to privacy in his home, while he does not have a right to privacy in his car.

However, the facts can be altered to make the legal rule more difficult to apply. Does an officer need a warrant to search a recreational vehicle (RV)? What if the inhabitant of the RV does not have a home and lives in the RV? What if that RV is parked on public property? What if the RV is parked in a Walmart parking lot? What if the RV is parked on the owner's private lot? What if the RV has its wheels replaced with concrete blocks? Creating an argument in these more ambiguous cases will require understanding not only the factual differences between the instant case and prior precedent, but also understanding the underlying legal principle for the rule.

AI can be used for same field searches by looking for cases that deal with search and seizure, focusing on similar facts such as "automobile," "recreational vehicle," or "home." AI would likely return legal precedents that have similar factual similarities to prior legal cases and the instant facts. These types of searches might be informed by supervised learning type models.

AI could also be used for same problem searches by looking for cases that have the same legal theories that underpin the legal rule. In this case, AI might be able to identify possible legal precedent that addresses the public policy concerns that underlie the search and seizure rules, even without specific labels surrounding those concerns. Hence, AI could identify the seemingly diverse connections between cases that may not be immediately evident to lawyers. For example, a model may return results where the cases discuss privacy and privacy-adjacent rights in different contexts, such as medical privacy rights or internet privacy rights. These types of searches might be informed by unsupervised learning type models.

#### B. AI AND LEGAL SEARCH

An application of supervised learning techniques to legal searching could be as simple as coding data from cases that have been recently decided. For example, cases involving illegal search and seizure could be used as inputs, with the judge or jury's decision regarding whether an illegal search was conducted labeled as the outcome. The supervised learning model could identify data such as the geographic location or item found within the cases and use those features as a way of making

predictions. Then, future cases with similar or identical facts could be identified by the machine and scored to determine the percent similarity to prior cases. Thus, using the cases found by the AI search, a lawyer could more easily identify the most relevant cases that have the best odds of supporting their legal argument.

In the legal context, supervised learning could help find and analyze the most relevant legal precedent. Supervised learning might help explain how the cited cases are similar or different from the instant facts. Supervised learning could even generate a score that indicates the similarity or dissimilarity between the case and the current factual situation. In this case, AI would be ordering the persuasiveness of arguments based on the factual similarities between the precedential case and the current factual situation. AI would be engaging in precedential reasoning, determining which legal precedents are similar enough to apply to the instant factual situations.

Another technique would be to have lawyers label just a few of the documents in a corpus, and then have the AI tool apply this style to the remaining documents. For example, the United States Patent and Trademark Office (USPTO) recently used an LLM to create the Artificial Intelligence Patent Dataset.<sup>27</sup> The LLM was trained on a few hundred patents and patent applications from eight AI component technologies covering areas such as natural language processing, AI hardware, and machine learning. The LLM then created a novel AI dataset by searching over 13.2 million patents and patent publications. A similar system could be used to efficiently classify legal opinions and legal statutes as well as secondary materials such as legislative history and law review articles.

Going back to the overarching search and seizure example, legal research based on the same field would search on cases dealing with keyword terms such as “homes,” “automobiles,” or “search and seizure.” The AI would search based on supervised learning to determine the cases with the most similar facts to help the lawyer find relevant cases. This type of search would be based on the ability of AI to sift through the vast amount of both useful and irrelevant information to return those cases, statutes,

---

27. Alexander V. Giczy et al., *Identifying Artificial Intelligence (AI) Invention: A Novel AI Patent Dataset*, 47 J. TECH. TRANSFER 476, 476–505 (2022).

and secondary sources that may be relevant to the specific facts at hand.

These types of searches would likely bring back many cases dealing with factual situations based on homes, automobiles, or search and seizures based on similar facts.<sup>28</sup> Accordingly, this type of search may simply be a small technological advancement over the current state of affairs. This type of search driven by supervised learning would be most useful when making precedential arguments and might be dispositive if a case with identical facts was found.

In this way, AI may simply act as a better version of a search tool, retrieving the most relevant cases and secondary sources from the most relevant jurisdictions. Accordingly, for the hypothetical, if AI could find multiple cases with identical facts and identical outcomes, the legal analysis step may be straightforward. Previous studies have shown that AI as a tool might be especially suited for this type of task. For example, Choi and Schwarcz demonstrate that students who utilized GPT-4 experienced a 29% increase in performance on a straightforward multiple-choice law exam.<sup>29</sup>

### C. LIMITATIONS OF AI AND LEGAL SEARCH

There are two main issues that may arise when using LLMs for legal search. First, use of a generative AI tool may potentially breach the confidential client information used to prompt the AI tool. Second, the AI tool may not be able to consider the client's unique circumstances, meaning the tool may not be able to grasp important legal nuances when proffering search results.

First, there are significant concerns regarding client confidentiality and use of AI tools. Lawyers have an obligation to keep client's confidences, and uploading a client's data in the form of a prompt to a third-party platform may breach that obligation. Developers of LLMs warn their users to not upload

---

28. Legal search providers are incorporating generative AI tools that are trained using self-supervised learning techniques to improve their search results. See *LexisNexis Launches Lexis+ AI, A Generative AI Solution with Linked Hallucination-Free Legal Citations*, LEXISNEXIS (Oct. 25, 2023), <https://www.lexisnexis.com/community/pressroom/b/news/posts/lexisnexis-launches-lexis-ai-a-generative-ai-solution-with-hallucination-free-linked-legal-citations>.

29. Jonathan H. Choi & Daniel Schwarcz, *AI Assistance in Legal Analysis: An Empirical Study*, 73 J. LEGAL ED. (forthcoming 2024).

sensitive information. For example, OpenAI's privacy policies, effective as of January 31, 2024, make clear that they collect certain data about their users and how they will use the data.<sup>30</sup> Under their policy, OpenAI may provide that data to third parties without notice to the user.<sup>31</sup> OpenAI also makes clear in its terms of use for its products that it may use its user's content, including their inputs and prompts, "to provide, maintain, develop, and improve" the company's services.<sup>32</sup> Additionally, OpenAI has a separate privacy policy for users who purchase "ChatGPT Enterprise" which is an AI tool built for businesses.<sup>33</sup> Under the Enterprise license, OpenAI makes clear that they "do not train on your business data[,] . . . you own your inputs (where allowed by law) . . . [and] you control how long your data is retained."<sup>34</sup>

Second, AI will only retrieve cases and legal materials based on the prompts given. However, it may be difficult to include the unique circumstances associated with each client. For example, in the hypothetical, a client may not want to reveal the fact that she is living in her automobile. This could bias certain outputs or pose challenges for the AI search tool in generating the most pertinent caselaw or legal information.

### III. LEGAL REASONING AND CREATING THE LEGAL ARGUMENT

Legal reasoning is a method of thought and argument used by lawyers and judges when applying legal rules to specific facts.<sup>35</sup> Legal reasoning requires the synthesis of legal rules, legal precedents, legal principles, public policy, and community values to analyze and solve legal problems. Legal reasoning requires the lawyer to think logically and systematically apply legal principles to the facts in order to reach a conclusion or create a legal argument. Legal reasoning is one of the most

---

30. *Privacy Policy*, OPENAI (Nov. 14, 2023), <https://openai.com/policies/privacy-policy>.

31. *Id.*

32. *Terms of Use*, OPENAI (Nov. 14, 2023) <https://openai.com/policies/terms-of-use>.

33. *Enterprise Privacy at OpenAI*, OPENAI, <https://openai.com/enterprise-privacy> (last visited May 3, 2024).

34. *Id.*

35. EDWARD H. LEVI, AN INTRODUCTION TO LEGAL REASONING 2 (2d ed. 2013).

difficult types of reasoning for students and practitioners alike, and it underpins the decision-making process within the legal system.

Legal reasoning often involves the use of precedent and analogy.<sup>36</sup> Precedential reasoning is where an earlier legal decision is applied because the two cases are essentially the same.<sup>37</sup> Analogic reasoning involves the application of an earlier legal decision, but the decision is not identical to the earlier one.<sup>38</sup> With analogic reasoning, lawyers must accent the similarities while discounting the differences between the earlier decision and the facts of the instant case.<sup>39</sup> Thus, after the legal searching is complete, lawyers then synthesize this information to form a legal opinion and develop strategies for the instant case.

Complicating the issue is the fact that legal reasoning may be based on legal principles that have no relationship to the facts of the instant case. For example, privacy cases that relate to internet privacy may be relevant to the analysis of search and seizure rules but have no relationship to the facts of the automobile or the home.

One challenge lawyers face when engaging in legal reasoning is determining whether two cases are similar enough to apply precedent (precedential reasoning) or justify extending the precedent through using analogy (analogical reasoning). By utilizing unsupervised learning techniques, a search algorithm might be able to retrieve related cases, not based only on factual similarities, but based on similar legal theories.

In legal reasoning, there is an inherent bias towards preserving the existing rules. Indeed, this default by courts to preserve the legal status quo has been termed *stare decisis*, Latin for “to stand by things decided.”<sup>40</sup> However, even in the

---

36. Precedent and analogy cover the majority of the types of legal reasoning lawyers engage in, including arguments based on rules, customs, and principles.

37. Grant Lamond, *Precedent and Analogy in Legal Reasoning*, STAN. ENCYCLOPEDIA PHIL. (Jun. 20, 2006), <https://plato.stanford.edu/entries/legal-reas-prec/>.

38. *Id.*

39. LEVI, *supra* note 35.

40. *Understanding Stare Decisis*, AM. BAR ASS'N (Dec. 16, 2022), [https://www.americanbar.org/groups/public\\_education/publications/preview\\_home/understand-stare-decisis/](https://www.americanbar.org/groups/public_education/publications/preview_home/understand-stare-decisis/).

face of this bias, the law is not static and allows for change when societal values change.

In 2001, Cass Sunstein asked the question “[c]an computers, or artificial intelligence, reason by analogy?”<sup>41</sup> His answer to this question was “they cannot[.]”<sup>42</sup> Sunstein came to this conclusion because in 2001, he believed that artificial intelligence was “unable to engage in the crucial task of identifying the normative principle that links or separates cases.”<sup>43</sup> The question of whether generative AI systems are capable of analogical reasoning, a hallmark of human intelligence, is debated, with some researchers concluding a tentative yes<sup>44</sup> and others a resounding no. Regardless of whether generative AI systems are at a level that can fairly be called reasoning by analogy, they are absolutely at a level where they are impacting lawyers’ ability to do so.

#### A. UNSUPERVISED LEARNING AND LEGAL REASONING

The central point behind analogical reasoning is that it is evaluative, and value driven. The difficulty associated with analogical reasoning lies in its reliance on societal principles, which are inherently influenced by societal values and can change over time. Additionally, analogical reasoning is difficult when the law is too ambiguous or when the rules are fragmentary, imprecise, or incomplete in describing the facts of the case. Accordingly, simply grouping factually similar cases may not be useful for analogical reasoning.

Going back to our overarching hypothetical, analogical reasoning requires the lawyer to think about the underlying principle for which the initial case was determined. In the case of our hypothetical, it is the principle that we have an expectation of privacy in our home that we do not have in our automobile. Accordingly, understanding how both the officer and the inhabitant view the vehicle and whether there was an expectation of privacy is the critical step needed for analogical

---

41. Cass R. Sunstein, Abstract, *Of Artificial Intelligence and Legal Reasoning* (U. Chi. Pub. L & legal Theory, Working Paper No. 18, 2001).

42. *Id.*

43. *Id.*

44. Taylor Webb, *Large-Scale AI Language Systems Display an Emergent Ability to Reason by Analogy*, 7 NATURE HUM. BEHAV. 1426 (2023); Taylor Webb et al., *Emergent Analogical Reasoning in Large Language Models*, 7 NATURE HUM. BEHAV. 1526 (2023).



reasoning. The expectation of privacy, however, is often rooted in values and is not strictly confined to the factual rubric of homes and automobiles.

Exacerbating the problem is the significance of perspective. For the hypothetical, the vehicle operator may believe that she has an expectation of privacy in her RV because she does not have a house but lives in her RV. However, the officer might not know that she lives in her vehicle. Thus, the perspective of the person applying the rule may have a dispositive effect on the outcome.

In this hypothetical, it is less useful (and likely impossible) to have a case for every single possible combination of facts. For example, a factual scenario with only five different variables would generate 120 unique cases. Thus, it is unlikely that every factual situation would be covered by a precedential opinion based on identical facts. Additionally, simply classifying and ranking the similarities and differences does not explain the societal values that created the rule. Understanding the principle behind the rule better helps to understand the outcome of a case based on specific facts. Thus, understanding the relevant similarities and differences is key.

#### B. USE OF AI FOR LEGAL REASONING

When the factual situation is complex, and if supervised learning techniques did not retrieve relevant results, then AI might help by searching using unsupervised learning techniques. This type of search would be based on the same problem. In the overarching search and seizure example, this may bring up information that may not deal with the same factual situation, but information based on the principles underlying the search and seizure analysis, such as privacy rights. For example, this type of search powered by unsupervised learning may bring results such as law review articles recognizing the societal value in information privacy, considerations of privacy rights within medical records, legal cases addressing information privacy within criminal proceedings, or instances involving internet privacy. None of these search results would have a direct bearing on the facts of the instant case but would help establish the privacy principles that justify the search and seizure rule. Employing this type of search could prove beneficial when constructing analogical

arguments, thereby aiding lawyers in advocating for changes in the common law grounded on evolving values or societal shifts.

AI tools that utilize unsupervised learning techniques may help lawyers find underlying principles between cases, thereby aiding them in engaging in analogical reasoning. Previous studies, however, indicate that AI still might not be the optimal tool for this particular task.<sup>45</sup> For example, the study by Choi and Schwarcz show that AI assistance provides less value for difficult issue-spotter questions.<sup>46</sup>

#### C. HARMS ASSOCIATED WITH USE OF AI FOR LEGAL REASONING

In a way, AI relies on precedent, because machine learning algorithms are necessarily constrained by their data training sets. If practitioners rely too heavily on AI to draft their legal arguments, they risk stifling legal change.<sup>47</sup> In this way, AI cannot be creative because it necessarily relies on the information that was fed into its algorithm. Overreliance on AI for creating legal arguments could result in a legal framework that fails to adapt to changing societal values. Instead, the law might become entrenched in the societal snapshot present at the time of the algorithm's creation. Legal precedent could solidify much faster and might not be easy to change, which could represent a significant loss to fast-changing areas of law.

Similarly, AI may fail to address emerging legal issues because its information corpus and algorithm may lack the specific linkage required to address these novel legal issues.

### IV. LEGAL WRITING

Legal writing is the culmination of a lawyer's legal search, legal reasoning, and legal arguments. It is a specialized form of communication to convey legal analysis, arguments, opinions, and information. Legal writing documents serve various purposes including persuading a factfinder, advising a client, and documenting legal transactions. Successful and effective legal writing is precise, clear, and adheres to established legal conventions. Legal writing can take many forms, including legal memos, briefs, contracts, legal opinions, articles, journals, and

---

45. Choi & Schwarcz, *supra* note 29.

46. *Id.*

47. See Amy Cyphert, Sam Perl & S. Sean Tu, *Artificial Intelligence Cannibalism and the Law*, 23 COLO. TECH. L.J. (forthcoming 2024).

judicial opinions. Effective legal writing requires an understanding of the law, the ability to analyze complex legal issues, and most importantly, the ability to articulate those concepts in a clear and persuasive manner.

#### A. AI AND LEGAL WRITING

AI will likely be an important tool used by practitioners. In fact, LexisNexis has already created an AI tool called Lexis+ AI.<sup>48</sup> This tool can help a lawyer draft and analyze legal documents. For example, given the correct prompts (including a prompt for the correct jurisdiction), this tool can create complex documents such as contracts, leases, and cease & desist letters. Interestingly, Lexis has two buttons where the user can “make this more aggressive” or “make this less aggressive” as well as a button that “explain what changed” between the previous and current version of the documents drafted.

AI tools will not only change the way that lawyers write.<sup>49</sup> Indeed, the study of how changes in technology impact writers has long intrigued scholars. “Whether clay, parchment, or screen, a writer’s materials affect word choice, prose, and style.”<sup>50</sup> The transition from typewriters to computer word processing made the process of revision much easier, and some argue that “the distinction between revision and composition began to erode entirely.”<sup>51</sup> The introduction of generative AI into lawyers’ writing toolkits will likewise change the way that lawyers conceive of arguments and structure them. Some of these changes may be welcome as they will increase lawyers’ efficiency and decrease their costs. However, there are other potential harms that are important to be aware of and guard against.

---

48. LEXISNEXIS, *supra* note 28.

49. See Bill Tomlinson et al., *ChatGPT and Works Scholarly: Best Practices and Legal Pitfalls in Writing with AI*, 76 SMU L. REV. F. (forthcoming 2024) (“If researchers and scholars rely too heavily on AI-assisted writing tools, they may lose their ability to write effectively without the assistance of these tools[.] . . . This atrophy could have a negative impact on the quality of scholarly writing and the ability of researchers and scholars to produce original and thought-provoking work.”).

50. Brad Desnoyer, *E-Memos 2.0: An Empirical Study of How Attorneys Write*, 25 J. LEGAL WRITING 213 (2021).

51. Matthew Kirschenbaum, *How Technology Has Changed the Way Authors Write*, NEW REPUBLIC (July 26, 2016), <https://newrepublic.com/article/135515/technology-changed-way-authors-write>.

## B. PRACTICAL HARMS TO LEGAL WRITING

Legal writing is a process, and the process itself may reveal new arguments or new perspectives that a lawyer did not initially consider. A lawyer who overly relies on generative AI may well undercut their own creativity and weaken their legal arguments, as the iterative process of writing is often crucial to the development of legal arguments in a brief. One of the authors has taught appellate advocacy, including appellate brief writing, for over a decade and has frequently noticed how students only realize after completing a first draft that they need to do additional research because their argument has several holes.<sup>52</sup> As one lawyer put it:

In the process of unwrapping the cold reality of our incompetence, writing also shows the way forward. It forces us along the journey that we must travel to reach that level of professional excellence that spells success. As we write we begin to see the holes in our logic, the need for more research, the need to use more precise language from authorities to support our arguments. Without writing, we can't do our job well even if we have a law clerk or associate who can produce the requisite written document for our case.<sup>53</sup>

These anecdotal reports are borne out by studies on how lawyers write. In an ethnographic study of the habits and practices of law firm associates in Pittsburgh, Pennsylvania, the researchers described how the associates would use the drafting process to really refine their ideas and arguments. As they were writing, the attorneys “engaged in brainstorming and planned out what they wanted to say[.] . . . They revised their plans as they read more or began to write, making more notes and annotating existing notes. As they tried to decide on a theory or thesis, they spent a lot of time reviewing and refining.”<sup>54</sup> This concept of writing as thinking and organizing is not unique to lawyers, as many other fields also recognize that the process of writing helps transform one’s understanding and learning.<sup>55</sup> Put

---

52. Professor Cyphert has taught Appellate Advocacy at the West Virginia University College of Law since 2011.

53. Anthony Renzo, *Don't Forget to Listen to What Your Writing Is Telling You*, 41 VT. B.J. 34 (2015).

54. Ann Sinsheimer & David J. Herring, *Lawyers at Work: A Study of the Reading, Writing, and Communication Practices of Legal Professionals*, 21 J. LEGAL WRITING INST. 63, 104 (2016).

55. See, e.g., Ian J. Quitadamo & Martha J. Kurtz, *Learning to Improve: Using Writing to Increase Critical Thinking Performance in General Education Biology*, 6 CBE LIFE SCI. EDUC., 140, 140 (2007) (concluding that students

another way, “[w]riting can change the writer, opening up new perspectives and beliefs or revealing what there still is to learn.”<sup>56</sup>

When lawyers skip the initial drafting phase of their writing, instead choosing to immediately use a generative AI tool, they potentially undercut their own creativity<sup>57</sup> and could end up with a written product that makes weaker legal arguments than they would have developed by doing the hard work of thrashing around in an initial draft.<sup>58</sup> The research suggests that in a straightforward, boilerplate legal analysis, using a generative AI tool for the first draft may not be that consequential. However, the more novel and complex the legal issue is, the more may be lost by using generative AI.<sup>59</sup> As Professor Harry Surden, an expert on AI and law, noted in a recent interview,

I think for certain basic, non-complicated legal cases, we’re not far from the day where a technology similar to GPT-4 can create a solid first draft of a motion that can, with significant double-checking and additional analysis, be ready to file. I think for more complicated cases that form the backbone of many law practices, these technologies should be treated as ‘first-draft’ machines rather than fully fledged motion-producing products.<sup>60</sup>

Lawyers should be careful even using generative AI as a first-draft machine and mindful that there are times that there can be real benefits to lawyers for going through the process of initial drafting the old-fashioned way.

---

whose laboratory experience included a writing assignment “significantly improved critical thinking skills whereas the nonwriting group did not”).

56. Rachele Garbarine, *How is AI Changing How We Write and Create?*, NC STATE UNIV.: HUMAN. & SOC. SCI. NEWS (Mar. 27, 2023), <https://chass.ncsu.edu/news/2023/03/27/how-is-ai-changing-how-we-write-and-create/>.

57. See Tomlinson et al., *supra* note 49, at 4–5 (“Some argue that the use of AI in scholarly writing undermines the integrity of the scholarly enterprise, as it may lead to a decrease in originality and creativity.”).

58. See Renzo, *supra* note 53.

59. See Choi & Schwarcz, *supra* note 29, at 6 (describing an empirical study of law students assigned to use GPT-4 and concluding that the findings “suggest that AI assistance might not be particularly useful on average in complex legal reasoning tasks (like essay-writing) that more closely resemble the difficult work of lawyering”).

60. Ronald M. Sandgrund, *Who Can Write a Better Brief: Chat AI or a Recent Law School Graduate?*, COLO. LAW., July/August 2023, at 28.

## CONCLUSIONS

The integration of AI in the legal field presents a transformative impact on legal search, legal reasoning, and legal writing. The capabilities of AI have significantly improved over the last few years and have the potential to revolutionize the efficiency and precision of legal professional in accessing relevant information, analyzing legal precedents, and generating insightful legal documents. The use of AI in legal research may facilitate identification of pertinent cases, statutes, and regulations, thereby streamlining the traditionally time-consuming and costly process of legal search.

Moreover, AI-driven legal searching could evolve to provide nuanced insights into complex legal queries, enabling practitioners to navigate intricate legal landscapes with greater accuracy. The technology's ability to comprehend context, recognize patterns, and adapt to evolving legal nuances positions it as a valuable tool in enhancing a lawyer's ability to conduct legal searches. For legal writing, AI has already proven itself as a potential tool for automating routine tasks, thus allowing legal professionals to focus on higher-order analytical and strategic thinking.

However, the promise of AI must be tempered with caution. AI hallucinations, alignment, privacy, confidentiality, and ethical issues, and structural biases will likely play a significant role in limiting the use of AI in the legal field. As lawyers navigate this evolving landscape, the synergistic relationship between AI and legal practitioners holds the potential to redefine the standards of legal research, legal search, and legal writing. AI tools will hopefully be used to create a more responsive, equitable, and effective legal system.

\*\*\*