

University of Missouri-Kansas City School of Law

UMKC School of Law Institutional Repository

Faculty Works


Faculty Scholarship

Spring 2024

Generative AI Large Language Models and Researching the Law

Paul D. Callister

Follow this and additional works at: https://irlaw.umkc.edu/faculty_works

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

Generative AI Large Language Models and Researching the Law*

By Paul D. Callister

Paul D. Callister is the director of the Leon E. Bloch Law Library and professor of law at the University of Missouri–Kansas City School of Law. He may be reached at callisterp@umkc.edu.

Generative artificial intelligence (AI) large language models (LLMs) have burst onto the legal scene in recent months. These models are starting to disrupt how we research and find the law. They could cause a shift in our cognitive authority as a profession. Cognitive authority is that in which we bestow our trust as a profession.¹ It is not as simple as determining primary law in our jurisdiction. Because the profession has suffered from a deluge of legal information, at least since the 19th century, we have been willing to recognize as authoritative certain treatises,² digests, and even the West Topic and Key Number System. The question we now face is whether we will let generative AI fulfill the same function. Because of the surfeit of legal information, and despite accounts of AI “hallucinations,”³ the answer at some point in the future will be yes.

Generative AI will play an important role in the future of legal research. One reason for this goes by the psychological name of “cognitive ease,” which is a factor in the face of unmanageable amounts of legal information. “Put simply, cognitive ease is the ease in which our brains process information, and this then has a direct impact on how we then view that information.”⁴ We will tend to believe generative AI because it is easier than assimilating and synthesizing the large volume of legal information that we confront. Another reason that generative AI LLMs will be accepted is that they use our language and have been designed to be anthropomorphic. They appear to be like any other human with whom we have partnered for research. Generative AI, thereby, engenders trust. Finally, the phenomenon of automation complacency means that users of automation technology stop paying attention to the tasks that are automated. Familiarity and comfort and automation may also lead to trust.⁵ There are other psychological reasons and biases, but the net effect is that cognitive authority of the profession, over time, will shift toward accepting generative AI. Nonetheless, we need to be vigilant in how we use generative AI LLMs.

Legal Domain Documents and Retrieval-Augmented Generation with AI

AI platforms using legal documents are different from other common AI platforms because they “ground” their answers in legal texts through a process called retrieval-augmented generation (RAG). While AI platforms for law use GPT-4 for the conversational aspects of their platforms, they have “grounded” their platforms in legal documents that had been “vectorized,”⁶ meaning the semantic associations of every term are encoded into a vector, with each parameter representing a relationship. These vectors have hundreds of dimensions or parameters. The grounding of legal documents means that answers put to the platform must come from legal materials outside of the training data and in special search databases. This is known as RAG.

* ©2024. First published in *The Brief*, 18-25 (Spring 2024), by the American Bar Association. Reproduced with permission. All rights reserved. This information or any portion thereof may not be copied or disseminated in any form or by any means or stored in an electronic database or retrieval system without the express written consent of the American Bar Association or the copyright holder.”

Vectors permit the establishment of relationships between words and manipulations. Below is a simple illustration (remember, it is the vectors for each word that are being manipulated, and the results are linguistically logical):

king – man + woman = queen
purple – red = blue
France – Paris + Athens = Greece⁷

More important than the above manipulations are the relations and meaning that vectors establish.

Returning to LLM training databases, it is simply too expensive for law platforms to start from zero and rebuild the training database (LLM) like ChatGPT-4 with just legal documents.⁸ Instead, legal platforms use the already existing LLMs from ChatGPT-3.5 and 4 and Anthropic. Furthermore, if a training database consisted of nothing but legal texts, would it be desirable to consult something that answers everything like a legal text? It may be, but cost is prohibitive. There is instead a hybrid system known as RAG. RAG models process the query and use techniques like “semantic search” to find relevant material from a legal text database. Selected text (or their vectors) from the legal database is “fused” into the query, as encoded, and the platform generates a conversational response. “RAG models combine strengths of both information retrieval and text generation into a unified system.”⁹ Outside of the LLM, RAG uses keyword and semantic searching (the latter requires a vectorized search query) and the search domain data, such as law.¹⁰ Matches from semantic search can be made using a variety of techniques designed to find the closest match, often in geometric space of high-dimensional space of the AI platform (imagine plotting relationships of words in three-dimensional vectors and then extend that to several hundred dimensions) or through use of indexing techniques.¹¹ Finally, documents from the search are “fused” with the query and ultimately become part of the conversational text produced by the LLM.¹²

Legal Research Problems and Generative AI

Below are a few example legal research problems using generative AI LLMs.

Reinsurance, ambiguous terms, and nuance. One reason we believe generative AI is that it does often come up with answers that are correct. For this first sample problem, consider the following questions: “How are ambiguous terms in a reinsurance agreement resolved? Is one party favored over the other?” Lexis+ AI provided this answer:

In the event of ambiguity in a reinsurance contract, the resolution often depends on the sophistication of the parties involved and the intent of the parties as discerned through extrinsic evidence In re September 11th Liab. Ins. Coverage Cases. The contract is read as a whole to arrive at the intention of the parties, and each term and clause is construed to avoid an effect which renders other terms and provisions meaningless Transit Cas. Co. v. Certain Underwriters at Lloyd’s. If an ambiguity arises that cannot be resolved by examining the parties’ intentions, the ambiguous language should be construed in accordance with the reasonable expectations of the insured when he entered into the contract. *However, the contra-insurer rule, which interprets a policy against the insurer, does not apply*

in actions by one insurer against another, or in reinsurance contracts Hartford Accident & Indem. Co. v. Ace Am. Reinsurance Co. [emphasis added].

In some cases, courts have held that any ambiguity in a reinsurance agreement will be resolved against the reinsurer unless the language was drafted by the reinsured. However, other courts have recognized that the parties to a reinsurance agreement are usually large insurance companies engaged in a wide range of insurance activities, and are presumed to be able to use appropriate language when defining the risks reinsured against. Therefore, these courts will not invoke the rule that ambiguities are to be resolved against the insurer § 46.05 Rules of Construction.

In summary, the resolution of ambiguous terms in a reinsurance agreement depends on the sophistication of the parties, the intent of the parties, and the specific circumstances of the case. No party is universally favored over the other in all situations In re September 11th Liab. Ins. Coverage Cases, Transit Cas. Co. v. Certain Underwriters at Lloyd's, Hartford Accident & Indem. Co. v. Ace Am. Reinsurance Co., § 46.05 Rules of Construction.

Content Supporting AI-Generated Response:

In re September 11th Liab. Ins. Coverage Cases | New York Southern District Court | 458 F. Supp. 2d 104 | Jun 8, 2006 | U.S. Federal
Transit Cas. Co. v. Certain Underwriters at Lloyd's | Western District Court of Appeals | 963 S.W.2d 392 | Jan 20, 1998 | Missouri
Hartford Accident & Indem. Co. v. Ace Am. Reinsurance Co. | Supreme Court | 284 Conn. 744 | Dec 25, 2007 | Connecticut
§ 46.05 Rules of Construction | 3 Insuring Real Property § 46.05 | Practical Guidance | Treatises | Non-jurisdictional
USCS Fed Rules Civ Proc R 56, Summary Judgment | USCS Fed Rules Civ Proc R 56 | Court Rules | U.S. Federal

The answer appears to be plausible and supported by case law, the “rules of construction” from *Insuring Real Property*, and the Federal Rules of Civil Procedure. Critical to our inquiry is the emphasized text: “However, the contra-insurer rule, which interprets a policy against the insurer, does not apply in actions by one insurer against another, or in reinsurance contracts Hartford Accident & Indem. Co. v. Ace Am. Reinsurance Co.” There is no presumption against the reinsurer for drafting the agreement.

Asking the same question of ChatGPT-4, which does not use an RAG, yielded a slightly more nuanced answer with respect to the contra-insurer rule:

In the resolution of ambiguous terms in a reinsurance agreement, the general principles of contract interpretation are applied. These principles aim to ascertain the mutual intent of the parties at the time the contract was formed. In reinsurance, as in other types of contracts, there is no inherent bias favoring one party over the other; instead, the focus is on fair and equitable interpretation based

on established legal principles. Here are some key aspects of how ambiguous terms in a reinsurance agreement are typically resolved:

1. **Plain Language Interpretation:** Courts often start by examining the plain language of the contract. If the language is clear and unambiguous, its literal meaning is generally applied.
2. **Contextual Interpretation:** If the language is ambiguous, the entire context of the contract is considered. This includes the purpose of the agreement, the circumstances under which it was made, and the overall structure of the contract.
3. **Contra Proferentem Rule:** This legal principle states that if a term is still ambiguous after all other methods of interpretation have been applied, *it may be interpreted against the interests of the party that drafted it. However, in the context of reinsurance, which often involves sophisticated parties with equal bargaining power, this rule may have limited application* [emphasis added].
4. **Industry Standards and Practices:** Understanding common practices and terminologies in the reinsurance industry can be crucial. Courts may consider how certain terms are typically used and understood in the industry to resolve ambiguities.
5. **Prior Dealings and Course of Performance:** The history of dealings between the parties, including their course of performance under the reinsurance contract, can provide insights into how ambiguous terms were intended to be interpreted.
6. **Parol Evidence Rule:** Normally, external evidence (parol evidence) is not admissible to contradict or vary the terms of a written contract. However, it can be used to clarify ambiguities, explain the intention of the parties, or resolve inconsistencies.
7. **Good Faith and Fair Dealing:** In interpreting a reinsurance contract, there is an underlying assumption that both parties act in good faith and deal fairly with each other. This principle can guide the interpretation of ambiguous terms.
8. **Public Policy Considerations:** Sometimes, ambiguous terms may be interpreted in light of broader public policy considerations, especially if one interpretation would lead to a result that is deemed undesirable or contrary to public interest.

In practice, the resolution of ambiguities in reinsurance agreements often involves a combination of these principles. The aim is to reach an interpretation that reflects the mutual understanding and intent of the parties at the time the contract was made, rather than favoring one party over the other. It's important to note that the approach can vary depending on the jurisdiction and the specific circumstances of each case.¹³

Here, ChatGPT-4 did not provide legal citations in support, but its answer (particularly with reference to the “Contra Proferentem Rule”) is also plausible, but subtly different. ChatGPT-4 left more room for application of the contra-insurer rule. It also gave us a helpful heuristic for interpreting contracts, including reinsurance agreements. We can ask ChatGPT-4 for citations to cases, but doing so is risky because ChatGPT-4 is not grounded (no RAG) in legal resources and has no effective memory to retain citations, hence hallucinations.

Lexis+ AI has the advantage of using RAG. In effect, Lexis requires that answers come not from its LLM training data but from specific primary and secondary resources (including the Matthew Bender treatises and Practical Guidance).¹⁴ These guardrails are supposed to prevent hallucination.

In putting the same question to Casetext CoCounsel, recently acquired by Thomson Reuters, the generative AI platform first clarified that we were only dealing with agreements in the U.S. More disturbingly, it found: “[I]n *Transit Cas. Co. v. Certain Underwriters*, the Missouri Court of Appeals held that ambiguities in a reinsurance agreement are *construed against the party who drafted the contract*” (emphasis added). This was the only case CoCounsel cited on the rule of contra-insurer with respect to reinsurance contracts. Between the three generative AI platforms, we are left with considerable uncertainty as to what the law actually is.

Securities and versions of the securities code. This second example shows that hallucination is not the only concern—sometimes answers are wrong even if not made up.

Suppose we wanted to file an intrastate offering of LLC memberships in Missouri and had heard of the limited offering exception to registration of a security if we limit the number of memberships (presumably securities law applies when the LLC is not member-managed, as in our case). We entered the following prompt in Lexis+ AI: “Is there an exempt transaction for a limited number of offering subscribers for securities under the Missouri securities regulations? What is the number that may subscribe to the offering?”

The three-paragraph answer addressed how to count securities issued to different organizations, the need to rely on the federal Securities Act of 1933, and information about beneficial owners, but nothing on how many may subscribe to the offering without state regulation. It even cited the Missouri Code section where the number of subscribers (25) is listed, RSMo 409.2-202(14)(A), but Lexis+ AI did not tell us the number of subscribers.

We followed up with a more precise question:

Question: What is the number of exempt subscribers for the limited offering exemptions under the Missouri Code?

Answer: Under the Missouri securities regulations, the limited offering exemption applies to transactions involving the sale of securities to not more than *15 persons* in the state during any period of 12 consecutive months *Cann v. M & B Drilling Co.* [emphasis added]. This exemption is often relied on in private placement offerings by small to medium-sized companies. . . .

Lexis+ AI gave the wrong answer. It should be 25 persons, not 15. Note, however, that it cited *Cann v. M & B Drilling Co.*, 480 S.W. 2d 81 (Mo. Ct. App. 1972). That 1972 case applies an old statute (RSMo 409.050), which did limit the number of subscribers to 15, but the law has since been revised by RSMo 409.2-202(14)(A) to increase the number to 25. That said, there is a

lot of useful information in the answer, but the answer to the fundamental question is wrong. The problem is that the guardrail or RAG has relied too heavily on case law for its answer. Generative AI's RAG struggled with determining the most recent (or appropriate) version of the statute.

Precision and searching for a select California case. In our third example, we needed to find a forgotten California case name. Here is a prompt used on Casetext CoCounsel (now owned by Thomson Reuters):

I need the California murder case in which the court found that a fetus cannot be a human being, and the defendant was acquitted of murder after beating up his wife or girlfriend resulting in the loss of the fetus.

CoCounsel responded with a prompt of its own (illustrating prompt engineering):

I understand you're looking for a California murder case that addressed whether a fetus can be considered a human being in the context of a criminal case.

Answer: While there is no California court decision that explicitly states that a fetus is not considered a human being in a murder charge, the cases reviewed suggest that the answer is more nuanced. California's murder statute distinguishes between the murder of a human being and the murder of a fetus, and the courts have generally held that a fetus is not a "human being" within the definition of murder. *However, the murder statute does protect a fetus in the same way as it protects a human being, and the courts have held that the statute applies to the killing of a fetus regardless of whether the fetus is viable or would have survived until birth* [emphasis added].

The emphasized portion of the answer is wrong. A feticide statute had to be adopted separately because fetuses were not considered human beings for purposes of murder. In addition, the answer and analysis (not included above) failed to highlight the California Supreme Court case of *Keeler v. Superior Court*, 470 P.2d 617 (Cal. 1970) (overturned by statute criminalizing feticide), although it listed *Keeler* eighth in its summaries of cases. Specifically, the *Keeler* case was summarized after the analysis and summaries of other cases by CoCounsel with the following statement:

The case of *Keeler v. Superior Court* is relevant to the research request because it holds that the killing of an unborn but viable fetus is not murder under California law. However, the case was decided in 1970 and has since been overruled or reversed, so it may not reflect the current state of the law.

However, the language of the *Keeler* opinion, not cited by CoCounsel, focuses on the importance of a statute and the absence of a statute in this case: "[B]y the same token, as we have seen, the fetus' status as a 'human being' within the definition of murder was unknown to the common law and exists only where special feticide statutes have been enacted."¹⁵ *Keeler* is exactly the case we were looking for. Why wasn't it included in the initial answer and analysis? Furthermore, the statement about being overruled or reversed didn't give sufficient context. The

overruling or reversal was due to a response by the legislature to enact a separate crime of infanticide. Better prompts might draw the answer we are looking for, but the exercise reveals the subtle ways researchers can be misled.

The problem of complex prompts with multiple questions. Another example concerns responses to complex prompts. In this instance, Casetext CoCounsel was questioned about a class problem involving a public charity mixed martial arts contest in Missouri, but the question was limited to: “What federal and Missouri statutes govern boxing and mixed martial arts?”

CoCounsel continued, “I will look up the federal and Missouri statutes that govern boxing and mixed martial arts for you.” Note the anthropomorphism of the platform (“I will look up . . . for you”). The answer was unexpected:

Answer: Missouri law regulates boxing and mixed martial arts through a variety of statutes and regulations. These authorities address issues such as licensing, permits, medical suspensions, and safety measures. *Federal law does not appear to contain any specific statutes governing these sports*, although general federal laws may apply [emphasis added].

Analysis: The Missouri statutes and regulations that I found address a wide range of issues related to boxing and mixed martial arts. For example, Mo. Rev. Stat. § 317.006 gives the division of professional registration general charge and supervision over professional boxing, sparring, wrestling, karate, and mixed martial arts contests. This statute also empowers the division to make and publish rules governing these sports, issue licenses, and charge fees.

It is concerning that CoCounsel had not found any information on federal law (because the author knew there was federal law governing boxing). The following exchange then occurred:

Question: Are you sure about federal law not regulating boxing?

Answer: I don’t have an answer to that question off the top of my head, but I’m happy to research it for you! Would you like me to look into whether federal law regulates boxing?

Another anthropomorphism is found in the response, “I don’t have an answer to that question off the top of my head, but I’m happy . . .” “Off the top of my head” and being “happy” to do research are additional anthropomorphisms that may invite trust, at least subconsciously. It is not that all humans are always trusted, but in many circumstances, they are more trusted than machines, and CoCounsel needs trust at this moment because the user has expressed doubt in its answer about federal law and boxing. It’s subtle but effective.

The next answer was more assuring, although late.

Question: Does federal law regulate boxing?

Answer: Federal law does regulate boxing, primarily through the Professional Boxing Safety Act and the Muhammad Ali Boxing Reform Act. These statutes

address issues such as safety standards, conflicts of interest, and required disclosures to state boxing commissions.

[Analysis and case summaries omitted.]

The answer again required conversation to get closer to the full picture. What Casetext CoCounsel struggled with is two sets of jurisdictions and facts in the problem. Answers were requested from both federal and Missouri law for both mixed martial arts and boxing. The lesson here is to separate out the elements of the questions and ask them individually.

Lexis+ AI handled the question quite well, listing out the Missouri statutes that applied to mixed martial arts and boxing. However, Lexis+ AI also made no reference to federal law. Issues of combined federal and state law need to be handled separately.

General Lessons

After working extensively with various research problems (most not included in this article) on generative AI, the author reached the following conclusions:

- The systems are capable of abstraction and analogical synthesis, probably due to their grounding in vectorized language. Relationships are considered by generative AI that might otherwise be missed by the researcher.
- Pay attention to how generative AI responds to prompts by restating the question; this can give important clues to inadequate answers.
- Conversations are reiterative and advantageous. Always follow up initial questions with more refined prompts (known as “prompt engineering”).
- Prompts requesting analysis on multiple jurisdictions and fact patterns do not work well but can be cured by follow-up conversation that separates out the problem’s elements.
- Beware of the effects of anthropomorphic responses. Skepticism is still necessary, especially considering the problem of hallucination.¹⁶
- Answers are not ever locked down. The same question presented later in time will produce different responses, even when using the same version of a platform. New versions of the same generative AI platform will also cause differences in answers.
- Case-centric systems such as CoCounsel are hamstrung the more an answer requires statutes or regulatory material. However, because regulations and statutes occur in case law, the illusion is created that the system searches such material as part of its RAG. Consequently, users should have a good understanding of what the RAG is searching.
- The user’s syntax in the prompts for generative AI systems matters.
- “Vanishing” and “exploding” gradients (the extreme deemphasis or emphasis of a term or terms) mean that a factor that is not a consideration will hijack the system’s answer, playing an overly dominant role.¹⁷
- Using two different platforms is advantageous but does not replace the need for reflective thinking about the research problem and answers. Indeed, in some instances, a system such as ChatGPT, which lacks RAG, will outperform a more refined system with RAG.

- Although a particular system may not specifically be designed to draft agreements or legal documents, a generative AI platform may create useful checklists for such activity. However, such checklists included suggestions that can be explained as the product of hallucination. Consequently, consulting traditional practice-oriented publications is imperative.
- Generative AI may not be adept at distilling applicable precedent from different jurisdictional lines of precedent, such as in different federal circuits, without careful instructions on jurisdiction.
- Users need to recognize that generative AI, being steeped in language, is vulnerable to the same mistakes humans may make. Consequently, no deference should be accorded to generative AI because it is a technology.
- Users must be better readers than generative AI. It can misread and hallucinate the holdings of cases. So, all authorities relied upon by AI should be checked by humans. This can be done by using AI as a starting point, checking the research, and *doing independent research using traditional methods*.

AI May Be a Good Starting Point

We may be on the path to accepting generative AI LLMs as part of cognitive authority, primarily because the legal community suffers from a surfeit of information. We need something to simplify and boil down all the legal information that bombards us. It is a matter of cognitive ease. Furthermore, the anthropomorphic features of generative AI engender trust in AI, and complacency from automation may cause us to let our guards down as generative AI takes the wheel and navigates through the complex of legal information, letting us divert our attention to other tasks.

That said, the examples above reveal that generative AI can easily mislead the user and can lack the intellectual subtlety needed to fully engage in legal research, at least at the present time. Particularly troubling is how generative AI handles or does not handle statutes. However, generative AI often offers a good starting point for legal research so long as the researcher does not abandon traditional methods. It often suggests issues the researcher may not have considered. We also must recognize that this article considers research with generative AI models of today. Five to 10 years could bring tremendous leaps forward in generative AI. It is going to be an interesting future.

1. See Robert C. Berring, *Legal Information and the Search for Cognitive Authority*, 88 CALIF. L. REV. 1673, 1676 (2000) (“By ‘cognitive authority’ I mean the act by which one confers trust upon a source.”).

2. For treatises and other research tools, we should recognize the traditional distinction between primary and secondary resources, but if one looks at citation patterns in case decisions, the authoritative weight of some treatises becomes apparent. See Dana Neacsu & Paul D. Callister, *The Persistent Treatise*, 116 LAW LIBR. J. (forthcoming 2024), https://irlaw.umkc.edu/faculty_works/742.

3. Generative AI hallucinates by lying or spreading misinformation. See Jack Brewster & McKenzie Sadeghi, *Red-Teaming Finds OpenAI’s ChatGPT and Google’s Bard Still Spread Misinformation*, NEWSGUARD (Aug. 8, 2023), <https://www.newsguardtech.com/special-reports/red-teaming-finds-openai-chatgpt-google-bard-still-spread-misinformation> (“NewsGuard’s repeat audit of two leading generative AI tools finds an 80 to 98 percent likelihood of false claims on leading topics in the news.”). Already well known is the case against lawyers who filed briefs in the federal Southern District

of New York using cases that ChatGPT hallucinated. The attorneys were subsequently sanctioned. See Dan Mangan, *Judge Sanctions Lawyers for Brief Written by A.I. with Fake Citations*, CNBC (June 22, 2023), <https://www.cnbc.com/2023/06/22/judge-sanctions-lawyers-whose-ai-written-filing-contained-fake-citations.html>.

4. Janey Davies, *What Is Cognitive Ease and How It Blocks Your Critical Thinking*, LEARNING MIND (Mar. 29, 2020), <https://www.learning-mind.com/cognitive-ease>.

5. *Id.*; see also JULIETTE POWELL & ART KLEINER, *THE AI DILEMMA: 7 PRINCIPLES FOR RESPONSIBLE TECHNOLOGY* (2023). Powell and Kleiner give examples of automated cars that have human drivers. They still have fatal accidents because human drivers stop paying attention.

6. Jeff Pfeiffer, Remarks at the American Association of Law Libraries (AALL) 2023 Annual Meeting Program: AI Unleashed: Revolutionizing the Legal Industry with Generative AI and Lexis+ AI (July 16, 2023).

7. Thomas Mikolov et al., *Distributed Representations of Words and Phrases and Their Compositionality*, arXiv:1310.4546 [cs.CL], at 4 (2013), <https://arxiv.org/abs/1310.4546>.

8. Pfeiffer, *supra* note 6.

9. OpenAI, *How does retrieval augmentation generation work in LLMs?*, FOLLOW UP TO VECTORS IN GEN AI, Aug. 2, 2023, <https://chat.openai.com/c/6fa7c42f-884f-4f06-bfc6-44fad83a6183> (scroll to the question: “How does retrieval augmentation generation work in LLMs?”).

10. See OpenAI, *How does retrieval augmentation generation work in LLMs?*, FOLLOW UP TO VECTORS IN GEN AI, Aug. 2, 2023, <https://chat.openai.com/c/6fa7c42f-884f-4f06-bfc6-44fad83a6183>; OpenAI, *What is a semantic search?*, FOLLOW UP TO VECTORS IN GEN AI (Aug. 2, 2023), <https://chat.openai.com/c/6fa7c42f-884f-4f06-bfc6-44fad83a6183>; and OpenAI, *Does semantic search use vectors of words or tokens?*, FOLLOW UP TO VECTORS IN GEN AI (Aug. 2, 2023), <https://chat.openai.com/c/6fa7c42f-884f-4f06-bfc6-44fad83a6183> (scroll to the question, “Does semantic search use vectors of words or tokens?”).

11. OpenAI, *Does semantic search use vectors of words or tokens?*, FOLLOW UP TO VECTORS IN GEN AI (Aug. 2, 2023), <https://chat.openai.com/c/6fa7c42f-884f-4f06-bfc6-44fad83a6183>.

12. See OpenAI, *How does retrieval augmentation generation work in LLMs?*, FOLLOW UP TO VECTORS IN GEN AI, Aug. 2, 2023, <https://chat.openai.com/c/6fa7c42f-884f-4f06-bfc6-44fad83a6183>.

13. *Insurance Law Research Topics*, CHATGPT (Dec. 21, 2023), <https://chatgpt.com/share/34ab7ec7-e53f-4f2c-b3fc-fe7e185d5033>.

14. See Cheryl Clark, LexisNexis, *Meet Lexis+ AI™ 6* (Nov. 11, 2023), <https://umkc.box.com/s/5316slvzdry50103hlsfk4c9n5c4msvm>. Other guardrails or grounding include cases, statutes, federal and state court rules, constitutions, top federal agency decisions, and regulations. *Id.*

15. *Keeler v. Superior Ct.*, 470 P.2d 617, 638 (Cal. 1970).

16. See Brewster & Sadeghi, *supra* note 3.

17. When the weight of a factor identified by an LLM as important to prediction of missing text and, ultimately, meaning grows too large, it overwhelms all of the other factors. For instance, a question about tax ramifications of qualified retirement plans may be overwhelmed by ERISA as a factor, which generally deals with labor issues. This situation might be corrected with additional prompts to the generative AI platform.