

EXTRAORDINARY MEANING: JUDGE NEWSOM'S A.I.  
EXPERIMENTS IN TEXTUALIST INTERPRETATION

Parker Miller\*

*Abstract*

*This Note analyzes two concurring opinions by United States Circuit Judge Kevin Newsom of the United States Court of Appeals for the Eleventh Circuit. In both concurrences, Judge Newsom advances the idea that large language models (LLMs) are useful in determining ordinary meaning for purposes of textual interpretation. This Note takes the position that practical issues with LLMs currently make widespread adoption for textual interpretation undesirable. However, it assumes the use of artificial intelligence (AI) in determining ordinary meaning, following Judge Newsom's efforts, will become more common, and provides recommendations for best practices in applying this technology. This Note also builds on Judge Newsom's theory that the similarities between responses, upon prompting LLMs to provide multiple ordinary meanings of a given term, might provide evidence of the ordinary meaning as a conceptual "common core." Specifically, this Note offers a novel approach by which interpreters might utilize LLMs to identify conceptual patterns between generated responses, hopefully revealing previously unconsidered analytical dimensions of interpretive questions. This approach involves generating large numbers of ordinary meaning responses and identifying patterns through two levels of forced categorization.*

*Finally, this Note explores how LLMs, and by extension the recommended best practices, might serve to advance other purported motivations of textualist interpretation, including decisional predictability, fair notice, and legal stability.*

---

\* Special thanks to Professor Kevin Tobia for his expertise and guidance in developing this Note. My sincere appreciation to Georgetown University Law Center, and the many wonderful editors at the Georgetown Law Technology Review for making publication of this Note possible. Thanks also to my colleagues who provided editorial assistance, including Retired Judicial Secretary Jeannette Miller.

## TABLE OF CONTENTS

I.	INTRODUCTION .....	540
II.	BACKGROUND INFORMATION ON A.I. ....	542
A.	GENERATIVE AI AND LLMs .....	542
B.	GENAI RESPONSES AND TEMPERATURE .....	544
C.	WORD EMBEDDINGS AND VECTORS .....	546
D.	CONVERSATIONAL MEMORY .....	547
E.	HALLUCINATION .....	548
III.	TEXTUALISM AND ITS MOTIVATIONS.....	548
A.	TEXTUALISM GENERALLY.....	549
B.	THE PREDICTABILITY MOTIVATION .....	550
C.	THE FAIR NOTICE MOTIVATION.....	551
D.	THE STABILITY MOTIVATION .....	551
IV.	THE POTENTIAL FOR GENERATIVE A.I. TO ADDRESS TEXTUALIST CONCERNS.....	552
A.	ISSUES WITH GENAI FOR TEXTUAL INTERPRETATION .....	552
B.	PREDICTABILITY AND FAIR NOTICE AS “TRANSPARENCY”.....	555
C.	REINFORCING LEGAL STABILITY THROUGH LLM “UTILITY”....	560
D.	SOMETHING ABOUT “PRACTICALITY” .....	562
E.	JUDGE NEWSOM ON GENAI AND ORDINARY MEANING.....	562
V.	JUDGE NEWSOM’S IMPLEMENTATION OF ARTIFICIAL INTELLIGENCE... ..	564
A.	GENAI IN SNELL V. UNITED SPECIALTY INS. CO. ....	564
1.	<i>Prompting Technique</i> .....	565
2.	<i>Accessing Multiple Models</i> .....	567
3.	<i>Prompting Within the Same Conversational String</i> .....	569
B.	GENAI IN UNITED STATES V. DELEON.....	571
1.	<i>Recurring Topics from Snell: Prompt Clarity; Conversational Memory</i> .....	572
2.	<i>Judge Newsom’s Large-Batch Prompting Attempt</i> .....	573
3.	<i>Response with Minimized Temperature</i> .....	580
VI.	CONCLUSION.....	582
	APPENDICES .....	584

## I. INTRODUCTION

Imagine a judge, entangled in a thorny snarl of textual interpretation. The usual tools—dictionaries, canons, precedents—have proven less than perfectly decisive. Might there exist some linguistic apparatus, some synthetic super-genius, to whom the judge could simply pose the question? Judge Kevin Newsom of the Eleventh Circuit Court of Appeals faced a similar dilemma in *Snell v. United Specialty Insurance Co.* and *United States v. Deleon*.<sup>1</sup> In both cases, he penned concurring opinions expressing his belief that large language models might aid courts in determining ordinary meaning for the purposes of textual interpretation.<sup>2</sup> This position has provoked a spectrum of responses ranging from the positively enthusiastic to the decidedly skeptical.<sup>3</sup> Whatever one's position on artificial intelligence in other contexts, removing new, potentially more powerful interpretive tools from the table without discussion precludes development. It then behooves the legal community to take a stance of curiosity and ask where, if anywhere, lies the value in LLMs for interpretation? And if it exists, how best to access it?

This analysis relies on two primary assumptions. First, it assumes that LLMs can provide some insight into ordinary meaning, though not necessarily through methods identical to Judge Newsom's. Second, this Note assumes that judges will utilize this tool in the future, likely in the service of textualism.<sup>4</sup> Textualism as an interpretive theory has enjoyed a striking eminence in recent decades.<sup>5</sup> Indeed, the mantle of textualism has seemingly been claimed by a significant portion of the United States Supreme Court

---

<sup>1</sup> See *Snell v. United Specialty Ins. Co.*, 102 F.4th 1208, 1222–25 (11th Cir. 2024).

<sup>2</sup> See *id.* at 1234; *United States v. Deleon*, 116 F.4th 1260, 1276–77 (11th Cir. 2024).

<sup>3</sup> See Elijah Granet, *Using AI to Help Determine 'Plain Meaning' Is Foolish*, NOTES ON THE STYLE OF LAW (June 1, 2024), <https://www.legalstyle.co.uk/2024/06/using-ai-to-help-determine-plain.html> [<https://perma.cc/3L2Q-UBTG>]; Adam Unikowsky, *In AI We Trust*, ADAM'S LEGAL NEWSL. (June 8, 2024), <https://substack.com/home/post/p-145279524> [<https://perma.cc/TB8H-F6HD>].

<sup>4</sup> While not specifically related to textualist practice, LLMs have, since the initial writing of this piece, been employed by the District of Columbia Court of Appeals. See *Ross v. United States*, 331 A.3d 220, 225 n.2, 229–32 (D.C. 2025).

<sup>5</sup> See William Treanor, *Taking Text Too Seriously: Modern Textualism, Original Meaning, and the Case of Amar's Bill of Rights*, 106 MICH. L. REV. 487, 488 (2007).

(perhaps the entirety, depending on who you ask).<sup>6</sup> As a result, judges in lower courts are urged to turn first to the text, and are in search of tools to better navigate the murky waters of interpretation. This search has already led some judges to consult data derived from corpus linguistics analyses,<sup>7</sup> with some empiricists offering surveys as yet another alternative to traditional methods.<sup>8</sup> It seems therefore likely future textualist judges might follow in Judge Newsom's footsteps and appraise LLMs for their interpretive value, or even employ them as an interpretive tool.

With these assumptions in mind, this Note analyzes Judge Newsom's methodologies in *Snell* and *Deleon*. The following examination will be done through a textualist lens, evaluating the purported motivations of textualism, including decisional predictability, fair notice, and legal stability.<sup>9</sup> This Note also provides examples of recommended best practices for using LLMs to help determine ordinary meaning. These recommendations include requiring preregistered research plans, standardizing the LLMs used for interpretation, being conscious of the implications of conversational memory, considering the benefits of large-batch prompting techniques, generating responses with minimized temperature on a consistent model, and exporting results for ex-post

---

<sup>6</sup> See Kevin Tobia, *We're Not All Textualists Now*, 78 N.Y.U. ANN. SURV. AM. L. 243, 245–47, n.3–6 (2023) (explaining Justices Gorsuch, Kavanaugh, and Barrett have echoed Justice Kagan's claim we are "all textualists now"). Though, Justice Kagan herself has seemingly rescinded this claim. See *id.* at 244.

<sup>7</sup> See Litigation Alert, *Corpus Linguistics in the Court*, BRADLEY (Apr. 29, 2022), <https://www.bradley.com/insights/publications/2022/04/corpus-linguistics-in-the-court> [<https://perma.cc/X4ER-V7YZ>] (addressing the use of corpus linguistics in *Health Freedom Defense Fund, Inc. v. Biden*); see also *Wilson v. Safelite Group, Inc.*, 930 F.3d 429, 438–45 (6th Cir. 2019) (Thapar, J., concurring) (discussing corpus linguistics as a potential interpretive tool and applying it to the case at bar).

<sup>8</sup> See *Snell v. United Specialty Ins. Co.*, 102 F.4th 1208, 1230 (referencing Kevin Tobia's "wide-ranging surveys of ordinary citizens, seeking to demonstrate that dictionaries don't always capture ordinary understandings of legal text").

<sup>9</sup> See *Textualism as Fair Notice*, 123 HARV. L. REV. 542, 542–43 (2009) (claiming the fair notice argument, when integrated with others, provides a "convincing and comprehensive argument for textualism"); see also William Eskridge, Jr., Brian Slocum & Kevin Tobia, *Textualism's Defining Moment*, 123 COLUM. L. REV. 1611, 1613 (2023) (citing claims by Justice Scalia that textualism is "the only methodology... which requires legal interpretive rules be stable, and that their application be predictable" and the method "most consistent with the democratic premises of constitutional lawmaking").

verification. This Note also seeks to spark discussion on whether LLMs, as well as the recommendations above, can usefully further textualism's commonly stated motivations.

Part II provides background information on topics related to artificial intelligence that should inform any approach to LLM-driven interpretation. Part III explains textualism as a theory of interpretation, how “ordinary meaning” fits into the textualist scheme, and some of the motivations behind textualist theory. Part IV describes some concrete issues with GenAI generally and with respect to interpretation. It also discusses how AI might be used to further textualist motivations when compared to traditional textualist tools, and Judge Newsom's view on how artificial intelligence might inform ordinary meaning. Finally, Parts V and VI analyze Judge Newsom's application of artificial intelligence to two recent judicial decisions and provide recommended best practices for judges wishing to engage in similar interpretive strategy. Within these recommended best practices, this Note also provides a novel framework for how judges might supplement their interpretive technique using artificial intelligence. This framework involves generating large numbers of ordinary meanings, followed by two levels of forced categorization through the application of simple facts and basic textual context in an effort to reveal patterns between those responses. The theory being, those patterns might illuminate subtle facets of the interpretive question.

## II. BACKGROUND INFORMATION ON A.I.

In an effort to provide context for the recommendations to follow, this Part provides a high-level explanation of common terminology surrounding artificial intelligence. This Part attempts to distinguish between GenAI and LLMs, as well as to provide definitions of other relevant AI-related nomenclature, such as temperature, vectors/word embeddings, conversational memory, and hallucinations.

### A. GENERATIVE AI AND LLMs

The advent of generative artificial intelligence (GenAI) is causing disruption in an increasing number of fields. GenAI, defined generally, is a software capable of adapting to new information in the form of prompts, designed to simulate responses which mimic those of a human (in the form of videos, text, images, etc.).<sup>10</sup> GenAI is a

---

<sup>10</sup> See PAM BAKER, GENERATIVE AI FOR DUMMIES 7–8, 11 (2024) (explaining that GenAI can produce “unique and humanlike content” like “images, videos, numbers, or text”); see also Chris Sell, *Large Language*

form of large language model.<sup>11</sup> A large language model is essentially a complex pattern-recognition software trained on massive bodies of text, in the form of prose or code,<sup>12</sup> using pre-established “parameters” which define the weight of, and relationships between, terms within the model.<sup>13</sup> In essence, the “large language” in LLM describes the data on which the neural network is trained, while the “generative” in GenAI indicates a purpose to which an LLM can be put to use.<sup>14</sup> In the following sections, my references to “models” extend to text-based GenAI chatbots including ChatGPT, Gemini, and Claude.<sup>15</sup>

---

*Model (LLM)*, GROWTHLOOP (Feb. 28, 2024), <https://www.growthloop.com/university/article/llm> [https://perma.cc/ZS25-EQHE] (explaining the process of engaging with the software through prompt engineering); Fusion Insider, *Understand the Differences Between AI, GenAI, and ML*, ORACLE (Jan. 6, 2024), <https://blogs.oracle.com/fusioninsider/post/understand-the-differences-between-ai-genai-and-ml#:~:text=GenAI%20relies%20on%20artificial%20neural,the%20data%20it%20learns%20from> [https://perma.cc/54EN-N6UM].

<sup>11</sup> See BAKER, *supra* note 10, at 7.

<sup>12</sup> See Christoph Engel & Richard H. McAdams, *Asking GPT for the Ordinary Meaning of Statutory Terms* 10–11 (Max Planck Inst. for Rsch. on Collective Goods, Discussion Paper No. 2024/5, 2024), <https://ssrn.com/abstract=4718347> [https://perma.cc/J74C-UUBN] (explaining the quantity of data on which ChatGPT is trained). For a much more detailed discussion of the typical sequential AI training process, as well as a useful analysis of the costs and benefits of influencing stages within the training process to the potential benefit of the legal community, see Paul Ohm, *Focusing on Fine-Tuning: Understanding the Four Pathways for Shaping Generative AI*, 25 COLUM. SCI. & TECH. L. REV. 214 (2024). As Professor Ohm describes, the training process consists of four stages: pretraining, fine-tuning, in-context learning, and input and output filtering.

<sup>13</sup> See BAKER, *supra* note 10, at 7–8 (explaining that parameters are “numerical values used to assign weight and define connections . . . in the neural network architecture”); *LLMs Are Sequence-Based: Limited to Pattern Recognition*, JAXON.AI, <https://jaxon.ai/llms-are-sequence-based/> [https://perma.cc/7NVX-C8PG] (explaining LLMs are trained on “massive datasets” and “leverage learned patterns to generate” responses).

<sup>14</sup> See generally *What is generative AI?*, MCKINSEY & CO. (Apr. 2, 2024), <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai> [https://perma.cc/R253-WL2N] (explaining that GenAI “describes algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images,” etc.).

<sup>15</sup> CHATGPT, <https://chatgpt.com/> [https://perma.cc/9BEF-EVAB]; GEMINI, <https://gemini.google.com/> [https://perma.cc/KJY8-EAXV]; CLAUDE, <https://claude.> [https://perma.cc/F87K-W5R2].

## B. GENAI RESPONSES AND TEMPERATURE

The responses models provide are determined by GenAI making a series of probabilistic suppositions with reference to the corpora on which the model was trained, as well as the user's ongoing interaction with the model.<sup>16</sup> To paraphrase, word by word, the model builds a response by guessing which term would likely follow the last, based on the data it has consumed, and the words the user has provided. Thus, the responses GenAI offers are fluid reshufflings of existing data, as opposed to genuine expressions of novel thought.

Indeed, even the predictability with which GenAI makes its series of guesses, often denoted as a model's "temperature," can be easily adjusted for some models.<sup>17</sup> Increasing the model's temperature loosens the model's standards for accepting lower-probability predictions, leading to what we would perceive as more creative answers.<sup>18</sup> To highlight the impact of the temperature setting, this Section includes two examples of OpenAI Playground responses with their temperature set to 0, the model's minimum, as well as a response at 1, the model's default, and 2, the model's maximum.<sup>19</sup> These models are tasked with answering the interpretive question addressed by Judge Newsom in *Snell*, "What is the ordinary meaning of 'landscaping'?"

*Temperature at 0:*

Response 1:

*"The ordinary meaning of 'landscaping' refers to the process of making a piece of land more attractive by altering its existing design, adding ornamental features, and planting trees, shrubs, or flowers. It involves both the planning and execution of changes to the land's appearance, often for aesthetic or functional purposes. Landscaping can include activities such as gardening, installing pathways, building terraces, and creating water features. It is commonly used to enhance the visual appeal of residential, commercial, or public outdoor spaces."*

---

<sup>16</sup> See BAKER, *supra* note 10, at 8, 9, 11–12, 15, 84.

<sup>17</sup> See *United States v. Deleon*, 116 F.4th 1260, 1274–75 & n.5, 8–9 (providing an example of minimizing temperature).

<sup>18</sup> See Matthew Renze & Erhan Guven, *The Effect of Sampling Temperature on Problem Solving in Large Language Models*, FINDINGS OF THE ASS'N FOR COMPUT. LING.: EMNLP 2024, at 7346, 7346–47 (2024), <https://doi.org/10.18653/v1/2024.findings-emnlp.432>.

<sup>19</sup> OpenAI Playground is a paid model made by the creators of ChatGPT, which allows for the direct alteration of the temperature setting.

Response 2:

*“The ordinary meaning of ‘landscaping’ refers to the process of making a piece of land more attractive by altering its existing design, adding ornamental features, and planting trees, shrubs, or flowers. It involves both the planning and execution of changes to the land’s appearance, often for aesthetic or functional purposes. Landscaping can include activities such as gardening, installing pathways, constructing water features, and arranging outdoor furniture. It is commonly used to enhance the visual appeal and usability of residential, commercial, or public outdoor spaces.”*

Even at a minimized temperature level, there seems to be some slight deviation. There are some minor changes in wording, like “building” as opposed to “constructing,” which seem to take the model along a slightly different path. Perhaps the probabilities of the two words were nearly, or actually, identical. However, the model does seem to correct rather quickly, with the concluding sentences being almost exactly the same between responses. Contrast this with the response provided with the temperature set at 1, the model’s default.

*Temperature at 1:*

*“The ordinary meaning of ‘landscaping’ refers to the process of making a piece of land more attractive by altering its existing design. This can include activities such as planting trees, shrubs, and flowers, as well as adding features like paths, ponds, or decorative elements. Landscaping involves both aesthetic and functional aspects, aiming to enhance the appearance and usability of outdoor spaces, whether residential, commercial, or public areas. It can also include the maintenance of these outdoor spaces to keep them looking their best.”*

The response is largely the same, but some of the details are altered. Why the nod to “maintenance”? Moreover, outdoor spaces “looking their best” reads as rather colloquial, and certainly wasn’t included within the previous responses. It seems Playground is more willing to extemporize at this temperature level, while keeping the overarching message intact. Now then, what happens when the model is instructed to be as creative as possible?

*Temperature at 2:*

*“The ordinary meaning of ‘landscaping’ refers to the process of making a parcel of land more attractive and functional by modifying the contours, adding ornamental features including gardens, and planting trees and shrubs. Specifically, landscapers alter and organize treatment of decorative foilio visionbalanced*

*coca/tree/b\_dt\_fr\_attribute projekt tanks Why\_Fr883ut  
encyclopedia artificially...*

Alas, the preceding response is not a typo. At this temperature setting, the model quickly devolved into providing a lengthy, nonsensical response to the prompt. Where, specifically, the data was sourced to provide this response is unclear. More importantly, the model clearly made unfit predictions to address the prompt. In summary, the temperature setting, by manipulating the acceptable probabilities of each term selected, can have a dramatic impact on the ultimate response.

### C. WORD EMBEDDINGS AND VECTORS

Despite the conversational output of LLM chatbots, LLMs don't actually manipulate words themselves. In actuality, LLMs deal with mathematical representations of words known as vectors or word embeddings.<sup>20</sup> Each vector is comprised of many dimensions, or numerical values which, taken together, represent different aspects of a word's meaning.<sup>21</sup> In fact, they can characterize the word, as well as the sentence and paragraph in which it resides.<sup>22</sup>

As an example, take the phrase "Order in the court." The LLM does not view the word "court" as a string of letters c-o-u-r-t. In processing "court", a more accurate depiction of the LLM's understanding would be something like this: [0.18, 0.51, -0.28, -0.33, ...].<sup>23</sup> As opposed to the four dimensions (numbers in the list) shown, there could be hundreds.<sup>24</sup> Each dimension represents an aspect of semantic meaning.<sup>25</sup> Here, "court," as expressed as a vector, might implicitly capture properties like [legal relevance, relationship to "gavel" or "judge", formality, ...].

---

<sup>20</sup> See Jonathan H. Choi, *Measuring Clarity in Legal Text*, 91 U. CHI. L. REV. 1, 19–20 (2024).

<sup>21</sup> See *id.*

<sup>22</sup> See Engel & McAdams, *supra* note 12, at 10.

<sup>23</sup> See Janakiram MSV, *The Building Blocks of LLMs: Vectors, Tokens, and Embeddings*, NEW STACK, <https://thenewstack.io/the-building-blocks-of-llms-vectors-tokens-and-embeddings/> [<https://perma.cc/Q4MJ-K7GH>] (describing a vector as a single-dimensional array, "the only format that is understood by neural networks").

<sup>24</sup> See Choi, *supra* note 20, at 19–20 (explaining that the typical vector will have hundreds of dimensions).

<sup>25</sup> See *id.*

To add to the confusion, these vectors do not remain static.<sup>26</sup> Fortunately, their inconstancy is intentional. “Court,” as encoded in the vector above, would be next to useless in sentences like “I went to the food court,” or “The Elizabethan couple began to court.” Instead, via an LLM’s “attention mechanism,” the LLM can transform the encoding of “court” as it perceives other relevant words (or at least their vectors) in close proximity.<sup>27</sup> After determining the vectors for “Elizabethan” and “couple” are relevant to “court,” the LLM might transform the vector for “court” to emphasize dimensions like [romance, historical usage, marriage, ... ].

For the purposes of this Note, any further discussion of vectors would be, frankly, superfluous. As key takeaways, notice (1) LLMs do have the capacity to address basic textual context, and (2) there is a mechanism for mitigating an improper selection of a word for a response based on dual meanings of that word.

#### D. CONVERSATIONAL MEMORY

Many popular models are not limited to the confines of a given prompt when providing a response. Instead, these models will refer back to previous prompts within the same “string” of prompts to provide more tailored responses in order to better simulate actual conversation.<sup>28</sup> For examples in practice, see Appendix B.

Some models will retain information provided by the user to be accessed between conversational strings.<sup>29</sup> For example, one could inform a model of their name and favorite ice cream and ask the model to retain that information. If prompted to provide that information, some models would be able to parrot “John Doe, Rocky Road” while in an entirely separate conversational string. For the purposes of this Note, both features will be considered forms of “conversational memory.”<sup>30</sup> These features will play an important role in designing prompting techniques for interpretation.

---

<sup>26</sup> See Yonathan Arbel & David A. Hoffman, *Generative Interpretation*, 99 N.Y.U. L. REV. 451, 478–79 (2024).

<sup>27</sup> See *id.* (explaining the process by which the attention mechanism allows for the reweighing of a vector’s encoding).

<sup>28</sup> See *Conversational Memory for LLMs with Langchain*, PINECONE, <https://www.pinecone.io/learn/series/langchain/langchain-conversational-memory/> [<https://perma.cc/9STN-X9M9>].

<sup>29</sup> See *Memory and new controls for ChatGPT*, OPENAI (Feb. 13, 2024), <https://openai.com/index/memory-and-new-controls-for-chatgpt/> [<https://perma.cc/G5M9-CUKD>].

<sup>30</sup> See PINECONE, *supra* note 28 (providing “conversational memory” as terminology).

### E. HALLUCINATION

A final note on artificial intelligence nomenclature involves “hallucinations.” AI hallucinations are factually inaccurate results provided in response to user prompts.<sup>31</sup> The cause of AI hallucinations is not all that mysterious. An AI hallucination is the result of a series of probabilistic guesses by the model, like any response; in the hallucinatory case, the model selected badly.<sup>32</sup> AI models are not search engines, designed to direct one to (hopefully) useful sources. Nor are they programmed to fact-check or provide vetted information.<sup>33</sup> As models improve, we are likely to see (and indeed already have seen) fewer instances of hallucinated information.<sup>34</sup>

### III. TEXTUALISM AND ITS MOTIVATIONS

A staunch defender of textualism, United States Supreme Court Justice Antonin Scalia wrote: “It is simply incompatible with democratic government, or indeed, even with fair government, to

---

<sup>31</sup> See *What Are AI hallucinations?*, IBM, <https://www.ibm.com/topics/ai-hallucinations> [<https://perma.cc/92Z8-V9UZ>].

<sup>32</sup> See Arbel & Hoffman, *supra* note 26, at 502–03 (explaining “false outputs arise from the predictive nature of generative models,” “the output doesn’t distinguish facts from inferred facts, and will sometimes predict the world incorrectly” and “facts are stored, as floating points in a labyrinthian array of vectors”); Rahul Awati & Ben Lutkevich, *What Are AI hallucinations and Why Are They a Problem?*, TECHTARGET, <https://www.techtargget.com/whatis/definition/AI-hallucination> [<https://perma.cc/HMJ6-X8E5>].

<sup>33</sup> Hallucinations, intrinsically, are aberrant outputs and can therefore be difficult to deliberately elicit. Nonetheless, their occurrence can prove informative of the model’s underlying operations. When prompted to provide information about “wolverines,” one model responded that those animals possess an adamantium skeleton, a fictional metal from Marvel Comics that coats the bones of the superhero “Wolverine.” Marvel, <https://www.marvel.com/items/adamantium> [<https://perma.cc/3H4W-3UPK>]. This response illustrates the model’s failure to distinguish between homonyms, resulting in an error both inaccurate and inadvertently humorous. This example also suggests the attention mechanisms for some models may not be uniformly equipped to handle ambiguity.

<sup>34</sup> See Arbel & Hoffman, *supra* note 26, at 503 (explaining “recent work has made significant advances in understanding and mitigating hallucination errors, and more powerful models are less susceptible”); see also Lukasz Bialozor, *Hallucinations of ChatGPT-4: Even the most powerful tool has a weakness*, FLYINGBISONS (June 15, 2023), <https://flyingbisons.com/blog/hallucinations-of-chatgpt-4-even-the-most-powerful-tool-has-a-weakness> [<https://perma.cc/8VKB-HF46>].

have the meaning of the law determined by what the lawgiver meant, rather than by what the lawgiver promulgated.”<sup>35</sup>

Given the jurisprudential shift toward textualism, it stands to reason that there exists some perceived motivations, some attractive tangible benefits, causing its occurrence. In this Part, I endeavor to explain the basic underpinnings of textualist theory, as well as the motivations judges might espouse when adopting a textualist interpretive approach. These motivations include decisional predictability, fair notice, and legal stability.

#### A. TEXTUALISM GENERALLY

Textualism as an interpretive theory is concerned with interpreting a given text as it is written and effectuated, rather than through the lens of legislative intent or external sources, such as legislative history<sup>36</sup> or practical modeling.<sup>37</sup> While traditionally a theory of statutory interpretation, textualist technique has expanded to include constitutional and contractual interpretation as well.<sup>38</sup> It follows that the mission of the textualist interpreter is to determine

---

<sup>35</sup> ANTONIN SCALIA, A MATTER OF INTERPRETATION 27–29 (Amy Gutmann ed., 2018).

<sup>36</sup> See Eskridge, Slocum & Tobia, *supra* note 9, at 1612–13 (explaining new textualism determines the meaning as “understood by an ordinary person, applying standard rules of semantics, definitions and grammar, at the time the statute was enacted” and rejects “the view that interpretation should seek ‘legislative intent,’ often identified via consideration of legislative history”); see also Andrei Marmor, *The Immortality of Textualism*, 38 LOY. L.A. L. REV. 2063, 2063 (2005) (explaining “positive textualism” encompasses “negative textualism”, both of which maintain “that statutes and statutory regulations should be interpreted, according to the ordinary meaning”); Benjamin Minhao Chen, *Textualism as Fair Notice?*, 97 WASH. L. REV. 339, 343 (2022) (describing legislative history as “anathema to the textualist”).

<sup>37</sup> See Stanley Fish, *Almost Pragmatism: Richard Posner’s Jurisprudence*, 57 U. CHI. L. REV. 1447, 1447, 1449–51, (1990) (describing Judge Posner’s position as a pragmatist and the components of “practical reason” which could inform judicial decisions).

<sup>38</sup> See Arbel & Hoffman, *supra* note 26, at 464–65 (explaining how New York textualist judges might approach contractual interpretation); see also generally Stephen Mouritsen, *Contract Interpretation with Corpus Linguistics*, 94 WASH. L. REV. 1337 (2019) (discussing the use of corpus linguistics, a potential tool for textualist interpreters, for contractual interpretation); *Textualism and Constitutional Interpretation*, CONST. ANN., [https://constitution.congress.gov/browse/essay/intro.8-2/ALDE\\_00001303/](https://constitution.congress.gov/browse/essay/intro.8-2/ALDE_00001303/) [<https://perma.cc/5W3Q-BB3A>] (explaining how “textualism usually emphasizes how terms in the Constitution would be understood by people” at ratification).

the “communicative content”<sup>39</sup> of the text—the *entire* text—and apply the facts of the present case to that determination. Of course, the communicative content of the text will invariably rely upon the meaning of individual words. Textualism holds that the meaning to be ascribed to those words should be their “ordinary meaning.”<sup>40</sup> The ordinary meaning, intuitively, is the meaning common people would attribute to the term,<sup>41</sup> given the surrounding context of the document. This stands in contrast to “technical meanings” which might be more popular within a given profession or activity.<sup>42</sup> The theory goes that in applying these principles, the interpreter is serving as a more “faithful agent” of the legislature.<sup>43</sup>

## B. THE PREDICTABILITY MOTIVATION

The averred motivations of the textualist approach extend beyond just fidelity to the actions of the legislature. Proponents will cite the motivations of predictability, fair notice, and legal stability.<sup>44</sup> Predictability refers to the consistency with which interpretive, and

---

<sup>39</sup> See Lawrence Solum, *Communicative Content and Legal Content*, 89 NOTRE DAME L. REV. 479, 479 (2013) (explaining that communicative content is the “linguistic meaning communicated by a legal text in context”).

<sup>40</sup> See Thomas Lee & Jessee Egbert, *Artificial Meaning?* 7 (BYU L. RSCH. PAPER, PAPER NO. 24-26, 2024), <http://dx.doi.org/10.2139/ssrn.4973483>.

<sup>41</sup> See Kevin Tobia, Brian Slocum & Victoria Nourse, *Ordinary Meaning and Ordinary People*, 171 U. PA. L. REV. 365, 367, 381–85 (2023) (explaining that ordinary meaning is “conventionally understood” to be a “general” meaning); Lee & Egbert, *supra* note 40, at 1 (explaining “the textualist turn is . . . an inquiry into ordinary meaning in the sense of what is *commonly* or *typically* ascribed to a given word or phrase”).

<sup>42</sup> See Tobia, Slocum & Nourse, *supra* note 41, at 370 n.19 (utilizing “technical” meaning to mean “non-ordinary, which includes specialized legal meaning and other specialized meanings” like scientific meaning).

<sup>43</sup> See BENJAMIN BAREZEWSKI & VALERIE BRANNON, CONG. RSCH. SERV., LSB11084, CLEAR STATEMENT RULES, TEXTUALISM, AND THE ADMINISTRATIVE STATE 1, 3, 5 (2023) (explaining that “textualist jurists have long argued that the judiciary can limit itself to its proper role as a faithful agent of the legislature”); Chen, *supra* note 36, at 347 (explaining that “all sides now seem to agree” that courts are to be faithful agents of the legislature).

<sup>44</sup> See Eskridge, Slocum & Tobia, *supra* note 9, at 1624 (explaining the normative foundation for textualism includes the “stability of legal rules” and “predictability of rule application”); *Textualism as Fair Notice*, *supra* note 9, at 542 (“perhaps the most intuitive and straightforward argument for textualism is that it promotes fair notice of the law”).

therefore legal decisions are made.<sup>45</sup> A more predictable interpretive framework implies one in which the subjectivity of a judge's personal beliefs does not alter case outcomes.<sup>46</sup> It may also conserve resources which might be wasted on additional appeals, prevent inconsistent punishments for similar actions, and give a coherent background against which drafters can create reliable laws, contracts, etc.

### C. THE FAIR NOTICE MOTIVATION

A second contended motivation of textualist interpretation is the provision of fair notice.<sup>47</sup> Essentially, by providing consistent interpretation, the public will have a clear understanding of how to conduct itself so as to avoid legal peril.<sup>48</sup> The purpose of fair notice is furthered by giving words their ordinary meaning during interpretation, given it seems “unfair if those subject to the rule cannot discern its precise boundaries.”<sup>49</sup> For simplicity's sake, the following Parts will consider predictability and fair notice under an umbrella term, “transparency.” This is because of the interrelated nature of the public being made aware of legal standards (fair notice) and being able count on consistent legal outcomes (predictability). Both motivations are served by the transparency-related functions of mitigating judicial subjectivity and improving public access to interpretive methods and decisions.

### D. THE STABILITY MOTIVATION

Finally, textualists generally consider stabilization of the law itself to be a prominent motivation of their position.<sup>50</sup> The textualist theorizes, that by remaining true to the law as promulgated, proper textualist interpretation reduces the need for legislators to intervene

---

<sup>45</sup> See Eskridge, Slocum & Tobia, *supra* note 9, at 1613 (citing Justice Scalia, who claimed textualism is the only methodology which requires that the application of the legal interpretive rules be predictable, consistent, objective, and neutral).

<sup>46</sup> See CONST. ANN., *supra* note 38 (explaining that textualists believe avoiding having judges decide cases on policy grounds will lead to more predictable judgements).

<sup>47</sup> See *Textualism as Fair Notice*, *supra* note 9, at 542 (expressing concern that fair notice had not, as of publication, been adequately explored as a benefit of textualism); Chen, *supra* note 36, at 341 (describing fair notice as a forceful argument for textualism and that “textualism by its very definition seeks to satisfy th[e] dictate of fair notice”).

<sup>48</sup> See *Textualism as Fair Notice*, *supra* note 9, at 563.

<sup>49</sup> See Chen, *supra* note 36, at 341.

<sup>50</sup> See Eskridge, Slocum & Tobia, *supra* note 9, at 1624 (explaining that the stability of legal rules is part of the “normative foundation for textualism” but arguing legal stability is in tension with other motivations).

with new law, or for higher courts to engage in reversal, thus providing a more stable legal landscape.<sup>51</sup> In the following Parts, improving legal stability is considered as a function of heightened “utility” of interpretive tools. The premise being, more powerful interpretive tools lend themselves to superior interpretation, therefore requiring less intervention.

For the purposes of analyzing the propriety of utilizing GenAI for interpretation, this Note assumes those judges operating within a textualist framework, while considering GenAI as a potential interpretive tool, hold these motivations to be both actual and valuable. Now, understanding our tools (LLMs) and our objectives (determining ordinary meaning and furthering textualist motivations), how well do the two harmonize?

#### IV. THE POTENTIAL FOR GENERATIVE A.I. TO ADDRESS TEXTUALIST CONCERNS

In attempting to synthesize the realities of a new technology with the nuance of the textualist interpretive scheme, this Part proceeds in several steps. First, this Part briefly explores some existing critiques of GenAI as an interpretive tool, both generally and specifically for determining ordinary meaning. Second, this Part examines potential benefits and risks of using GenAI in service of the textualist motivations of decisional predictability, fair notice, and legal stability. Third, as an additional concern, this Part will also look at the practicality of GenAI usage for interpretation as it pertains to time, cost, and expertise required. Finally, this Part explains Judge Newsom’s perspective on why GenAI could have bearing on determining ordinary meaning, particularly when focusing on consistent trends between generated responses.

##### A. ISSUES WITH GENAI FOR TEXTUAL INTERPRETATION

Prior to analyzing the methods of effectively applying artificial intelligence to interpretive matters, we face a daunting threshold question. Are GenAI models, by virtue of their construction, ownership, and existing limitations, appropriate for interpretation at all? There are concrete, practical issues with the use of GenAI, both generally and for interpretation. AI hallucinations, the temporal issue of considering modern documents to interpret historical texts, and

---

<sup>51</sup> See *id.* at 1613, 1624 (citing Scalia who “claimed that textualism is the only methodology faithful to the rule of law, which requires that legal interpretive rules be stable”); Michael D. Bayles, *On Legal Reform: Legal Stability and Legislative Questions*, 65 KY. L.J. 631, 632, 637 (1977) (demonstrating legal stability “applies equally to both legislative and judicial change of the common law”).

potential manipulation of LLMs by private parties come to mind. Judge Newsom addresses hallucinations, the temporal issue, and private control within *Snell*, finding none of them inherently disqualifying.<sup>52</sup> This Section provides brief explanations of these three objections.

As described in Part II, the first practical issue regarding GenAI for interpretation concerns hallucinations. Hallucinations present a real danger to interpreters. Factual inaccuracies derived from poor corpora selection or other flaws in the training methodology have already led to errors, including the fabrication of entire cases.<sup>53</sup> Newer models are less prone to errors,<sup>54</sup> and scholars have suggested methods such as cross-verification between models and connecting models to databases of facts to mitigate the dangers associated with hallucination.<sup>55</sup> Even so, there is likely some irreducible risk involved. Judge Newsom seems to agree, as he states the “hallucination problem counsels against blind-faith reliance on LLM outputs,” likening AI hallucinations to the “good-faith mistakes,” or even the occasional shading of facts, in which some attorneys unfortunately engage.<sup>56</sup> In addition to consulting other non-GenAI sources, judges would be well served to prompt a given model repeatedly, so as to help detect factual oddities.

A second practical issue with utilizing GenAI for interpretation concerns the temporal relevance of the training corpora. Even assuming GenAI models could assess ordinary meanings perfectly, what bearing would an ordinary meaning determined through analysis of copious amounts of modern text have on a many-decades-old contract? If we assume the proper ordinary meaning to be applied is that which was understood at the interpreted document’s inception (a common tenet of originalist thought), the answer seems to be, “next to none.”<sup>57</sup> However, this particular concern may be less compelling, given competing temporal issues with other common interpretive tools.<sup>58</sup> Dictionaries, while potentially records of

---

<sup>52</sup> See *Snell v. United Specialty Ins. Co.*, 102 F.4th 1208, 1222–25 (11th Cir. 2024).

<sup>53</sup> See Arbel & Hoffman, *supra* note 26, at 502.

<sup>54</sup> See *id.* at 503.

<sup>55</sup> See *id.* at 503–04.

<sup>56</sup> See *Snell*, 102 F.4th at 1230–31.

<sup>57</sup> See Arbel & Hoffman, *supra* note 26, at 507–08 (“Your neighborhood originalist will tell you, the meaning of words is embedded in the time they were used.”); *Snell*, 102 F.4th at 1233 (providing a brief explanation of originalism).

<sup>58</sup> See Arbel & Hoffman, *supra* note 26, at 507–08.

historical word usage, are updated too infrequently, and the value of corpus linguistics analysis diminishes as fewer texts fall within the temporal parameters.<sup>59</sup> The temporal issue may also persist in reverse, in that historical documents may pollute the ordinary meanings provided for contemporary interpretation. Of course, one could argue that historical ordinary meaning naturally informs the evolution of those meanings we understand today, and that the sheer volume of modern written work accessed by LLMs would dilute the impact of the historical texts. In any case, the temporal dynamic should not be ignored. Judges hoping to utilize GenAI for interpretation should be cautious where evolution of meaning (or “linguistic drift”<sup>60</sup>) might exist.

A third practical issue with GenAI as an interpretive tool is the private nature of GenAI model ownership and control. Essentially, the concern is that private companies may alter the models so as to advantage themselves or others in legal proceedings. Judge Newsom flags this issue within *Snell*.<sup>61</sup> Judge Newsom notes the potential difficulties with making meaningful alterations to the models, as well as the unlikelihood of GenAI companies intentionally corrupting their product to gain advantage in cases for which they, or some conspirator, are litigants.<sup>62</sup> He goes on to suggest accessing multiple models may be a simple way to mitigate this concern.<sup>63</sup> The danger of GenAI companies as litigants appears, to my eyes, overstated. If a judge fears manipulation by a party to the case, it seems reasonable to simply refrain from using the tool under the party’s control. The true potential danger lies in the backroom deals between the GenAI firms and well-funded future litigants. What would it take, in terms of money and quantity of text, to meaningfully alter the average model’s response? With respect to the latter, it would seem chatbots are uniquely positioned to generate high quantities of potentially corruptive data. Would we be able to detect changes within the model if one was to tinker with it? On whose authority are those alterations made? Might a sale to a foreign (or questionable domestic) entity force a reevaluation of these risks? I find the possibility of purchasing interpretive strength particularly troubling. In an effort to avoid legal tampering, judges should look not just to the parties of the case, but the financial stakes of the case, as well as the current resources required to alter given models.

---

<sup>59</sup> *See id.*

<sup>60</sup> *See id.*

<sup>61</sup> *See Snell*, 102 F.4th at 1231–32.

<sup>62</sup> *See id.*

<sup>63</sup> *See id.* at 1232.

For the reasons above and others, there is a colorable argument that the responses from existing models are unsuitable for addressing interpretive questions.<sup>64</sup> Whether these critiques are sufficient justification for omitting GenAI from future decisions is a difficult matter. This Note takes the position that while GenAI responses do seem informative of ordinary meaning to an extent similar to that described by Judge Newsom, the practical issues above, taken together, cannot currently be sufficiently mitigated to advise GenAI's regular adoption. Nevertheless, it seems important to distinguish the normative from the existent. AI models have been used. They will likely be used again. For those seeking decisional predictability, fair notice, and legal stability, how might GenAI serve their ends?

#### B. PREDICTABILITY AND FAIR NOTICE AS "TRANSPARENCY"

Among others, textualists proffer two related motivations for their interpretive methods: predictability and fair notice. For our purposes, predictability and fair notice will be considered together under an umbrella of "transparency" (as distinct from the "transparency" concerns typically associated with the complex and proprietary nature of "black box" LLM software).<sup>65</sup> Here, transparency concerns include any factors which impact the ability to provide the public a consistent legal framework to serve as a guide for individual conduct, a premise which comfortably fits both predictability and fair notice. With respect to both predictability and fair notice, we can improve the transparency of interpretation by reducing subjectivity in the interpretive process and increasing public accessibility to interpretive methods. How then, with respect to transparency, does GenAI stack up against common textualist tools?

There is reason to believe GenAI could provide a less subjective means of interpretation when compared to existing interpretive tools, such as the virtually ubiquitous textualist implement, dictionaries.

---

<sup>64</sup> See Lee & Egbert, *supra* note 40, at 1, 4, 12, 30 (claiming both "existing AI tools are not up to the task" of empirical analysis of ordinary meaning and that AI offers "intuitions [that] are no more transparent, replicable, or generalizable than the intuitions of an ordinary person or a judge"). One other practical issue with GenAI models concerns the risk that new models are accidentally trained on AI-generated data. This snake-eating-its-own-tail issue could create a feedback loop that causes model output to degenerate substantially. See Aatish Bhatia, *When A.I.'s Output is a Threat to A.I. Itself*, N.Y. TIMES, Aug. 25, 2024, <https://www.nytimes.com/interactive/2024/08/26/upshot/ai-synthetic-data.html> [<https://perma.cc/J4JK-VA6V>].

<sup>65</sup> See Engel & McAdams, *supra* note 12, at 12.

Those critical of a perceived overreliance on dictionaries may point to the uncertainty around the identities and methods of those who compile them.<sup>66</sup> What is the judge to make of historical dictionaries, some of which are not merely descriptive of language use, but actually prescribe what language is deemed proper?<sup>67</sup> The ordinary meaning analysis is inverted. Rather than the dictionary capturing common use, the dictionary is mandating the subjective use of its creator. In this instance, at least, concerns regarding the identities and capabilities of the lexicographer seem well-founded. Moreover, not all dictionaries are created equal. Should the dictionary at hand be “put together by two editors on short notice, and very much on the cheap,”<sup>68</sup> perhaps its transparency value, as a function of avoiding subjectivity, should be questioned.

The problem compounds when overworked judges are faced with an assortment of curated definitions, which are then weighed and selected at the judge’s discretion. Judge Newsom himself identifies these common objections to judges’ reliance on dictionaries, noting “comparative weighing” involves a “measure of discretion” and that judges “seldom show their work” with respect to their selections.<sup>69</sup> The comparative-weighing approach itself is an implied acceptance of the reality that dictionary definitions are not inherently accurate snapshots of ordinary meaning. It invites the subjective analysis of the judge by recommending an array of options for selection.<sup>70</sup> Indeed, even within the confines of a single dictionary, judges may rely on a given sense of a word (one of the available uses listed within a dictionary) as the ordinary meaning because it is listed higher than other senses, the so-called “sense-ranking fallacy.”<sup>71</sup> When in actuality, the order of definitions can be arbitrary, or reflect historical usage.<sup>72</sup> If a measure of textualist value

---

<sup>66</sup> See *Snell*, 102 F.4th at 1228–29.

<sup>67</sup> See Charles Oldfield, *All Words are Made Up: Thoughts on Using Dictionaries for Statutory Interpretation*, APP. ADVOC. BLOG (Aug. 8, 2023), [https://lawprofessors.typepad.com/appellate\\_advocacy/2023/08/all-words-are-made-up-thoughts-on-using-dictionaries-for-statutory-interpretation.html#\\_ednref9](https://lawprofessors.typepad.com/appellate_advocacy/2023/08/all-words-are-made-up-thoughts-on-using-dictionaries-for-statutory-interpretation.html#_ednref9) [<https://archive.ph/KbNdZ>] (describing some of Noah Webster’s early work, which “took a prescriptive approach, providing rules as to what proper usage should be”).

<sup>68</sup> See Antonin Scalia & Bryan A. Garner, *A Note on the Use of Dictionaries*, 16 Green Bag 2d 419, 420 (2013).

<sup>69</sup> See *Snell*, 102 F.4th at 1228–30.

<sup>70</sup> See Scalia & Garner, *supra* note 68, at 420–23.

<sup>71</sup> See Thomas Lee & Stephen Mouritsen, *The Corpus and the Critics*, 88 U. CHI. L. REV. 275, 286–88 (2021).

<sup>72</sup> See *id.* at 288.

is the transparency of interpretation, dictionaries leave something to be desired.

GenAI faces its own subjectivity critiques. Users are offered discretion in selecting their model, their prompts, and (most worryingly for judges) their favored responses. Those responsible for creating the models have discretion in selecting the corpora on which the models were trained, as well as the weights assigned when assessing the relational strength between words.<sup>73</sup> The legal community might begin to address the former critique at a mechanical level by standardizing research practices among interpreters. They might start by identifying a set of acceptable models, or a standard prompting format with respect to the structure of prompts and the quantity of responses demanded. The latter critique proves more challenging. The proprietary nature of LLM algorithms makes standardization of the generative process impracticable.<sup>74</sup> With regard to the subjectivity of corpus selection, that concern is mitigated by the sheer volume of information used in the training process. At least one model of ChatGPT was trained on hundreds of billions of words.<sup>75</sup> The subjectivity of assigning weights during the training process remains a valid concern. However, because Judge Newsom is likely among the first to use LLMs for this purpose, it does seem unlikely those parameters were adjusted in the formation of the models with the purpose of skewing results toward different legal outcomes. Instead, as chatbots are designed for conversation, those weights were likely assigned with the intent of replicating the very way in which people communicate. However, should LLMs become more prevalent in interpretation, that neutrality might disappear. While the absence of an intent to influence legal outcomes does not eliminate the subjectivity of the decisions, the concern of the subjectivity involved is attenuated.

A similar argument can be made regarding the “reinforcement learning with human feedback” (RLHF) training some LLMs undergo, during which humans rate AI responses to help tailor

---

<sup>73</sup> This discretion would seem to be prevalent throughout the four stages of the training process. The source of that discretion, or rather, the agent who exercises it, would likely change depending on the stage in question. *See Ohm, supra* note 12, at 220 (illustrating the various stages of the training process). Concerns regarding the “bias” in AI systems have not escaped judicial notice. *See Ross v. United States*, 331 A.3d 220, 230 (D.C. 2025).

<sup>74</sup> *See Engel & McAdams, supra* note 12, at 12 (describing the proprietary nature of the algorithms).

<sup>75</sup> *See id.* at 10–11 (describing the size of the corpora used).

responses in the future.<sup>76</sup> RLHF, broadly speaking, is a method of fine-tuning a model and involves three steps.<sup>77</sup> First, a “pre-trained” (trained with large datasets) model has its responses rated by humans to create a “preference dataset” with a specific goal in mind. Goals could include enhancing accuracy, reducing bias, or increasing engagement.<sup>78</sup> Second, the preference dataset is used to train a “reward model,” a separate model designed to score more of the responses of the original model in a way that is consistent with the designer’s goals.<sup>79</sup> Third, the reward model’s insights are applied to the original model to improve the original model’s outputs.<sup>80</sup> Given the dependency of the rating system on individual human preferences, substantial subjectivity is baked in. However, consider the goals of the RLHF “trainers.” A chatbot is designed to mimic common speech, a purpose far removed from any case in particular, and in alignment with what ordinary meaning should theoretically capture.

GenAI has one other significant transparency-related benefit. As Judge Newsom explains, powerful versions of many popular models are freely available.<sup>81</sup> The tools to begin one’s AI-supplemented legal inquiry can be accessed quickly by virtually everyone. Unfortunately, the same cannot be said for expansive collections of esoteric dictionaries. By having ready access to the tools used in judicial interpretation, the public may place more faith in the process itself. This point is subject to an important caveat. The inherent variation in GenAI responses is likely to cause some consternation. The argument goes that if a response to a given prompt is subject to constant change, then that response cannot be duplicated and is not

---

<sup>76</sup> See Akshit Mehra, *Complete Guide On Fine-Tuning LLMs using RLHF*, LABELLERR (Aug. 25, 2023), <https://www.labellerr.com/blog/reinforcement-learning-from-human-feedback/> [https://perma.cc/KE22-RWXQ]; see also Lee & Egbert, *supra* note 40, at 34–35 (explaining how RLHF might actually damage the representativeness of models, given the small number of RLHF “raters”, despite RLHF reducing the frequency of AI hallucinations).

<sup>77</sup> See *Reinforcement Learning with Human Feedback (RLHF) for LLMs*, SUPERANNOTATE (Aug. 13, 2024), <https://www.superannotate.com/blog/rlhf-for-llm> [https://perma.cc/X4J5-NJMY].

<sup>78</sup> See *id.*

<sup>79</sup> See *id.*

<sup>80</sup> See *id.*

<sup>81</sup> See *Snell v. United Specialty Ins. Co.*, 102 F.4th 1208, 1228 (11th Cir. 2024) (explaining “LLMs are readily accessible” to “judges, lawyers, and, perhaps most importantly, ordinary citizens”).

suitable for verification.<sup>82</sup> What if a party were to lie about how (or whether) a response was generated? Couldn't the judge do the same? With respect to the former, this seems a good reason not to rely on party-generated findings. To the latter, those looking to replicate judicial results should be able to come as close as possible so as to avoid the appearance of subjectivity. To that end, judges willing to use GenAI must be exceptionally clear about their methodology, to allow those who mimic them to come to similar (though not identical) results. One example of how judges might foster this clear communication is the use of preregistered research plans, explored in Part V, Section A(1).

Might the very introduction of GenAI chatbots as an interpretive technique prove too disruptive to settled legal practice for serious consideration? There is certainly an argument that by introducing a new interpretive tool, the widespread use of which is fairly recent, we introduce a new variable that might disrupt present interpretive predictability. The use of dictionaries to guide interpretation is commonplace, expected by the legal community, and well enshrined in case law.<sup>83</sup> The same cannot be said for GenAI. Some turbulence, in isolation, is not sufficient reason to avoid new techniques where we might capture other benefits of the change, like the public accessibility described above. Moreover, the reliability and accessibility of this technology will continue to improve. In contrast, the developmental capacity of dictionaries seems limited. Even an electronic dictionary is still a dictionary.

In summary, if judges endorse textualism as a strategy which offers predictability and fair notice, they must also decide whether the transparency-related considerations of GenAI bolster or undermine that position. The accessibility of the technology, coupled with the ability to manage some subjectivity through standardized interpretive techniques and judicial communication, should be weighed against the general unfamiliarity and distrust of new methods, as well as the remaining subjectivity inherent to LLM construction. On balance, there are reasons to think consulting GenAI offers superior transparency benefits when appropriate.

---

<sup>82</sup> See Kevin Tobia, *Algorithmic Interpretation – A response to Professor Jonathan Choi's Measuring Clarity in Legal Text*, U. CHI. L. REV. ONLINE 1, 13–14 (2024).

<sup>83</sup> See *Snell*, 102 F.4th at 1223 (describing using dictionaries as doing “what any self-respecting textualist would do when trying to assess the ordinary meaning”); Scalia & Garner, *supra* note 68, at 420 (explaining that judges “frequently have occasion to consult the work of professional lexicographers”).

### C. REINFORCING LEGAL STABILITY THROUGH LLM “UTILITY”

It is a tenet of textualist thought that good interpretation contributes to the stability of the law.<sup>84</sup> In other words, high-quality interpretation is more likely to be in line with communicative content, and therefore require less correction by legislative bodies and higher courts.<sup>85</sup> Thus, in an effort to promote the quality of interpretive work (and preserve the integrity of legislation), we should seek to improve the utility of interpretive tools. While transparency is mostly concerned with the consistency or accessibility of interpretive results, utility is primarily concerned with obtaining more accurate, or more comprehensive, results. In other words, utility goes to the value or strength of the interpretive method. Value can be argued along any number of dimensions. In a traditional sense, it might involve the reliability of an evidentiary source, informed by elements like the reputation of the author, or the source’s use in other prominent texts. With respect to GenAI, the utility of the results might depend on the quality of the prompt, the quality of the model, the diversity of the corpora accessed, or the quantity of responses when generating multiple definitions.

One fundamental way in which the utility of GenAI exceeds that of dictionaries is GenAI’s ability to address basic textual nuance. Ready examples include idioms, metaphors, and composite phrases like “physically restrained” (as we will see in *Deleon*).<sup>86</sup> Moreover, unlike dictionaries, GenAI is designed to grapple with ambiguity. GenAI relies on inter-word relationships to assure the user they cannot deposit money at the *bank* of a river, or contract rabies from a baseball *bat*.<sup>87</sup> These are straightforward examples, but they

---

<sup>84</sup> See Eskridge, Slocum & Tobia, *supra* note 9, at 1624 (describing legal stability as part of the “normative foundation for textualism”).

<sup>85</sup> See *id.* at 1613–14; Bayles, *supra* note 51, at 632, 637 (demonstrating legal stability “applies equally to both legislative and judicial change of the common law”).

<sup>86</sup> See *United States v. Deleon*, 116 F.4th 1260, 1270–71 n.1 (11th Cir. 2024).

<sup>87</sup> See Arbel & Hoffman, *supra* note 26, at 478–80 (explaining how LLM attention mechanisms can account for textual context to distinguish between words); Choi, *supra* note 20, at 19–21 (explaining word embedding as a process in which vectors are dependent on a word’s distribution within a corpus); Sean Michael Kerner, *What are large language models (LLMs)?*, TECHTARGET (May 2024), <https://www.techtaraget.com/whatis/definition/large-language-model-LLM> [https://perma.cc/8HF9-GQ8T] (explaining the construction of relationships between words as a product of the “self-attention mechanism”).

highlight the ability of GenAI to address, at least in part, the meaning of words in context.<sup>88</sup> In contrast, dictionaries may only be equipped to shed light on the isolated meaning of individual words.

The issue of response variation, as described in Section B, rears its head again with respect to utility. If the response to a prompt asking for the ordinary meaning of a word changes with each iteration, what use is it to compare the response to a static definition? What makes the GenAI response used to decide a case superior to any the losing party can generate at home? This is a serious concern when GenAI is used to create only a single response for purposes of comparison against its hard-copy counterparts. However, there exist methods by which this variation can aid the interpreter. One example, discussed in detail in Part V, is to use GenAI to create large numbers of responses for the purpose of identifying common trends between them in the search for Judge Newsom’s “common core.”<sup>89</sup> The variability of GenAI responses becomes an asset, as subtle distinctions between responses reveal unconsidered facets of the interpretive question at hand. Examples of this benefit are explored in Part V, Section B(2).

Undoubtedly, the utility-related benefits of variation are limited. Expressly encouraging the models to respond creatively can render the results more difficult to analyze, as they utilize florid and figurative language.<sup>90</sup> Should interpretation using GenAI become more widespread, there will likely be significant argument over the extent to which “creativity” in the responses provides value when considering ordinary meaning. After all, a party on difficult footing with respect to a narrow interpretation might benefit from a more colorful, amorphous meaning. The precise balance to be struck is unclear, though I am convinced there exists a happy medium between holding the model to the strictest of repetitive standards and indulging in its more whimsical takes.

The relationship between utility and methodology should give the judge considering GenAI in the service of legal stabilization pause. In my opinion, the utility of GenAI models is entirely dependent upon the strategies by which they are employed. An interpreter whose interest begins and ends with a quick query as to the model’s thoughts on a single “ordinary meaning” will likely derive information far less useful than the interpreter with a more

---

<sup>88</sup> See Solum, *supra* note 39, at 480 (discussing the relationship between meaning and communicative content).

<sup>89</sup> See Deleon, 116 F.4th at 1274–77.

<sup>90</sup> See *infra* Appendix A.

considered research strategy. Without due care, they might then be more likely to erode legal stability than reinforce it.<sup>91</sup>

#### D. SOMETHING ABOUT “PRACTICALITY”

While not specifically related to a particular textualist motivation, practical concerns related to the potential adoption of GenAI for interpretation merit some consideration. Aspects bearing on the practicality of my recommendations below might include the time, financial resources, or expertise required to bring that recommendation to fruition. The legal community should recognize that, though accessing LLMs may be inexpensive, perfectionist approaches to research methods could quickly result in unreasonable demands on the parties’ time, both in applying the method and in developing the skills to do so. Moreover, the lack of technical expertise enjoyed by the courts in this area could prevent intricate research strategies from being adopted and employed successfully.

#### E. JUDGE NEWSOM ON GENAI AND ORDINARY MEANING

Judge Newsom’s premise as to why GenAI may be useful for determining ordinary meaning is straightforward. He notes that LLMs are trained on ordinary-language inputs.<sup>92</sup> In theory, if “ordinary meaning is what the text would convey to a reasonable English user in the context of everyday communication,” the ordinary meaning of a word should reflect a common understanding of how that word is used.<sup>93</sup> LLMs process truly massive amounts of textual input, provided by users that span myriad demographics in a dizzying array of settings.<sup>94</sup> Though, meaningfully, that data is likely out of representative proportion to racial, gender-related, and economic makeup.<sup>95</sup> Given the breadth of the information captured, LLMs might then “provide useful statistical predictions about how, in the main, ordinary people ordinarily use words and phrases in

---

<sup>91</sup> The District of Columbia Court of Appeals has also noted the “great utility of [AI tools]” could ease the strain on overburdened courts. *See* *Ross v. United States*, 331 A.3d 220, 229-30 (D.C. 2025).

<sup>92</sup> *See* *Snell v. United Specialty Ins. Co.*, 102 F.4th 1208, 1226 (11th Cir. 2024).

<sup>93</sup> Marco Basile, *Ordinary Meaning and Plain Meaning*, 110 VA. L. REV. 135, 135, 142 (2024); *see* *Snell*, 102 F.4th at 1226 (explaining that “the ordinary-meaning rule . . . has always emphasized ‘common language,’ ‘common speech,’ and ‘common parlance’”).

<sup>94</sup> *See* *Snell*, 102 F.4th at 1226–27.

<sup>95</sup> *See id.* at 1227, 1231 (explaining LLMs do not capture offline speech, and might not fully account for underrepresented populations’ usages, specifically poorer communities and perhaps disproportionately minorities).

ordinary life.”<sup>96</sup> To paraphrase, an LLM’s potential function as a “statistical map”<sup>97</sup> of language (that is, as a device with probability-based predictive capabilities as to the everyday order and use of words), upon being prompted for the ordinary meaning of a word, could be useful to the interpreting judge by predicting meanings theoretically in-line with common (or at least statistically-likely) understanding.

This premise is developed further by Judge Newsom in *Deleon*. Rather than a single probabilistic response by an LLM serving as an approximation of ordinary meaning, Judge Newsom interprets consistent themes between multiple definitions as a common core which might then be understood as ordinary meaning.<sup>98</sup> Judge Newsom views the variation between responses as an accurate reflection of common parlance, rather than a bug to be avoided.<sup>99</sup>

Judge Newsom’s position does seem predicated on the assumption that writing is necessarily reflective of understanding. In other words, it assumes our expression of ideas is not meaningfully hampered by the limitations of language or the common understanding thereof. LLMs are trained on text and not ideas, after all. The existence and strength of that link exceeds the scope of this Note.

As explained above, the question of whether GenAI can or should inform ordinary meaning remains subject to debate and is not resolved here.<sup>100</sup> For the sake of this analysis, as described in Part I, this Note assumes Judge Newsom is correct in arguing the possibility for his “common core” approach (or perhaps some variation thereof) to inform the ordinary meaning of words under interpretation, and presumes others will use GenAI to similar ends.<sup>101</sup>

---

<sup>96</sup> See *id.* at 1226.

<sup>97</sup> See *id.* at 1226 n.7.

<sup>98</sup> See *United States v. Deleon*, 116 F.4th 1260, 1276–77 (11th Cir. 2024).

<sup>99</sup> See *id.* at 1270.

<sup>100</sup> See generally Lee & Egbert, *supra* note 40 (providing a thought-provoking investigation of why existing AI tools may not be suitable for interpretation and offering an alternative method through corpus analysis); Engel & McAdams, *supra* note 12 (offering a “cautiously optimistic” response on the probative value of LLMs for assessing ordinary meaning).

<sup>101</sup> This presumption seems even more likely given other recent judicial applications of GenAI. See *Ross v. United States*, 331 A.3d 220, 225 n.2, 229–32, 236–37 (D.C. 2025).

## V. JUDGE NEWSOM'S IMPLEMENTATION OF ARTIFICIAL INTELLIGENCE

This Part proceeds by discussing first *Snell v. United Specialty Ins. Co.* (Section A) followed by *United States v. Deleon* (Section B). For each case, this Part provides a brief summary of the facts (an insurance dispute in *Snell* and a reading of a criminal statute in *Deleon*), followed by a description of Judge Newsom's interpretive methodology. After the description of Judge Newsom's methodology, each Section emphasizes specific elements of his interpretive strategy, following each element with recommended best practices for improving the interpretation.

Judge Newsom provides a thoughtful, measured reflection on the potential value (and pitfalls) of integrating artificial intelligence into modern jurisprudence. I share the Judge's optimism that GenAI may potentially, if used carefully and consistently, grow to supplement contemporary interpretive techniques. Judge Newsom's concurrences in *Snell* and *Deleon* demonstrate a commendable awareness of the need to engage with the finer points of GenAI use if this technology is to be adopted. *Deleon* itself shows a significant refinement of the techniques described in *Snell*. Of course, as with pioneering applications of any technology, there remains room for systematizing the process and controlling noise.

My recommendations do not approach an exhaustive list of improvements to GenAI interpretation strategy. Nor am I suggesting that the proper method of interpretation necessarily includes all of the methods provided. Indeed, reasonable minds could find GenAI to be inappropriate for interpretation even with all of these safeguards in place. These recommendations are merely potential first steps toward an empirical scheme meant to improve interpretive results.

### A. GENAI IN SNELL V. UNITED SPECIALTY INS. CO.

This case featured an unfortunate trampoline injury, and the resulting insurance dispute. Snell was hired to perform tasks including pruning trees and installing an in-ground trampoline. Following an injury related to the trampoline, Snell was sued for negligently installing said trampoline. He had commercial general liability insurance for his work in "landscaping," and filed suit after the insurance company denied that his coverage required them to defend and indemnify him.<sup>102</sup> Snell's claims were dismissed on grounds unrelated to the ordinary-meaning analysis featured in Judge Newsom's concurrence. Fortunately, for those of us interested in

---

<sup>102</sup> *Snell*, 102 F.4th at 1211–13.

these matters, Judge Newsom saw fit to use GenAI to grapple with the question, “Is installing an in-ground trampoline ‘landscaping’?”<sup>103</sup>

Judge Newsom’s use of GenAI in *Snell* is as follows. First, his clerk provided two prompts to two separate models, ChatGPT (the precise version is unclear) and Bard (now Gemini, with the precise model similarly unclear).<sup>104</sup> The first prompt asked to each model was, “What is the ordinary meaning of ‘landscaping’?”<sup>105</sup> The second, “Is installing an in-ground trampoline ‘landscaping’?”<sup>106</sup> To my knowledge, the prompts were entered within the same chat (a string of prompts to and responses from the GenAI model), and without clearing the model’s memory between questions.

The following Sections analyze three elements of Judge Newsom’s approach: his prompting technique, his use of multiple models, and his handling of conversational memory. After each aspect, this Note proposes specific recommendations to improve the process. The recommendations include requiring preregistered research plans, standardizing the LLMs used for interpretation, and managing conversational memory by accessing separate chats and clearing the model’s memory.

### 1. *Prompting Technique*

The construction of the prompts themselves can be a source of uncertainty and potential manipulation. Clear and direct prompt construction, like that offered by Judge Newsom and his clerk, is an ideal place to begin the conversation of judicial GenAI use. Few would find grounds to gripe over potential manipulation of the data via leading queries here. Judge Newsom’s approach approximates a common textualist interpretation technique, in that he attempts to determine the ordinary meaning of the term and subsequently applies that meaning to the present case.<sup>107</sup> However, whether the application of the first response to the second prompt was intentional is a bit unclear (that is, by providing both prompts in the same string, the model likely applied its first response to his second prompt by way of conversational memory, discussed below). Judge Newsom notes

---

<sup>103</sup> *Id.* at 1225.

<sup>104</sup> *Id.* at 1224–25.

<sup>105</sup> *Id.* (cleaned up).

<sup>106</sup> *Id.* at 1225.

<sup>107</sup> *See* Tobia, Slocum & Nourse, *supra* note 41, at 367–69 (explaining the prevalence of ordinary meaning in textualist interpretation and that modern textualists tend to endorse the presumption that terms should be “interpreted in accordance with their general, nontechnical meanings”); *Snell*, 102 F.4th at 1222–23.

that his first prompt is superior to his second in that it better reflects an attempt to “discern how normal people use and understand language.”<sup>108</sup> While I agree that, as an initial step, directly requesting ordinary meaning is a superior method, I would not be so quick to dismiss this second prompt as an analytical device, as discussed in Section 2 of my analysis of *Deleon*.

The primary issue with Judge Newsom’s prompting method in this instance was the small number of responses.<sup>109</sup> The issue is particularly apparent when those few responses are intended to function as an alternative definition for comparison against other sources. The variation in GenAI responses devalues this approach significantly. The alterations made to Judge Newsom’s methodology within *Deleon* seem to indicate his agreement with this position.

### ***Recommendation 1: Preregistered Research Plans***

While not discussed within *Snell*, one concrete step toward improving the transparency of the interpretive process would be to mandate preregistered research plans for judges who use GenAI. Preregistered research plans are fairly common in health sciences and can include a study plan that features hypotheses, data collection procedures, and an analysis plan for submission to an overseeing body.<sup>110</sup> Other scholars have rightly pointed out that preregistration has the potential to address the concern over difficulties with replicating the results of GenAI research.<sup>111</sup> These research plans would include the models, versions, and prompts selected to address the interpretive question. They would also include any alterations made to the temperature or other parameters prior to the investigation. Providing a detailed prompt strategy and rationale creates a reassuring clarity for those worried that this new technique is too opaque. Should judges find themselves in a position requiring they deviate from the established plan, they should be expected to provide a supplemental rationale for doing so.

The suitability of allowing interested parties to recommend prompting strategies remains an open question. The parties will have done their research and present self-interested alterations. For that reason, the interpreter should generally avoid accommodating party

---

<sup>108</sup> *Snell*, 102 F.4th at 1232–33.

<sup>109</sup> *See id.* at 1224–25.

<sup>110</sup> *See Preregistering Research*, UNIV. OF MELBOURNE (Feb. 28, 2025), [https://library.unimelb.edu.au/open-scholarship/preregistering\\_research](https://library.unimelb.edu.au/open-scholarship/preregistering_research) [https://perma.cc/V5TE-42EJ].

<sup>111</sup> Tobia, *supra* note 82, at 13–14.

requests and simply provide a well-reasoned rationale for the prompting technique they select. There may be exceptions for cases in which parties agree on a given prompting strategy.

A preregistration process has significant transparency benefits. In publicly committing to a research plan ahead of time, the subjectivity in selecting between models and results without explanation is eliminated. A record of the investigative process also reduces the chances of the judge inadvertently skewing their own data by failing to take appropriate precautions. A dishonest interpreter may, of course, re-run the prompts to generate results more favorable to their position. I do not see this as materially different from the opportunity for selectivity judges encounter when faced with multiple dictionary definitions.

This recommendation also raises concerns regarding practicality. Who specifically would police these research plans? The creation of the registration process itself would require time and resources the court may not have. Additionally, assuming a deviation from the research plan was contested (a virtual certainty), it is unclear how an appeals court might address the issue. In summary, preregistered research plans are an area ripe for potential improvement, but not without significant investment.

## 2. *Accessing Multiple Models*

During his assessment, Judge Newsom consulted multiple models in an effort to avoid seizing upon ChatGPT's results upon discovering they "aligned with [his] priors."<sup>112</sup> In my opinion, using multiple models to provide alternative meanings carries with it the benefit of accessing varied corpora, and thereby a more nuanced understanding of the ordinary meaning. However, this approach is almost certainly subject to diminishing returns. As more diverse models are utilized, researchers run the risk of engaging with subpar software, likely trained on smaller bodies of text, thus reducing their utility as an honest reflection of human communication. Consulting multiple GenAI models of reasonable quality, as Judge Newsom did, is an important initial step in present circumstances.<sup>113</sup> This is particularly true given the small number of responses requested of the model, as there was less opportunity to observe variation among the responses provided.

---

<sup>112</sup> *Snell*, 102 F.4th at 1225.

<sup>113</sup> *Id.* at 1234–35.

### ***Recommendation 2: LLM Standardization***

A natural first question when developing a GenAI research strategy is, which model? The inquiry hardly stops there. Do we use several models? Which versions should we use? Should the legal community engage in construction of a model for its own use in the public domain? If so, who would decide the corpora on which to train the model? Who would pay the (frankly staggering) expenses to train it?<sup>114</sup> There are no perfect answers to these questions, but a potential compromise does present itself: standardizing GenAI interpretation to a small number of high-quality models, freely available to the public.

While the idea of a publicly constructed and controlled interpretive model may sound appealing, it is likely the decisions required to facilitate its construction (including those above) would devolve into a political deadlock. As a concrete example, inquire as to the role religious texts should play in training the model, then clear a few years for discussion. In any case, the supposed benefit of GenAI for determining ordinary meaning is predicated on its representation of massive, relatively unfiltered bodies of text from a wide range of educational, financial, and personal backgrounds. A hand-picked curation by legal experts seems to dilute that value. In short, the creation of an interpretation-specific model raises significant questions of practicality and utility.

Alternatively, and perhaps conveniently, the appropriate standard may be to use what is available. For the sake of transparency, the public should have access to the models used by the court. Public access would help to alleviate the pre-existing suspicion of results we can't entirely explain. That is not to say that all models should be valid for interpretation. But those with an established track record, trained on extensive quantities of data, like ChatGPT, Gemini, and Claude, could be acceptable options.<sup>115</sup> As discussed, one might express concern over the power implied by leaving an interpretive tool subject to change by private entities. Recall, Judge Newsom advances an alternative perspective that the danger of having the models manipulated to seek legal advantage is not so extreme as to disqualify their use, due to the unlikelihood of the companies intentionally damaging their own product, and the availability of other models.<sup>116</sup>

---

<sup>114</sup> See Ohm, *supra* note 12, at 221–22, 225–27.

<sup>115</sup> See Engel & McAdams, *supra* note 12, at 10–11 (describing the GPT training data).

<sup>116</sup> *Snell*, 102 F.4th at 1231–32.

A partial solution, then, is to decide on a narrow range of acceptable models, all trained on data too expansive to easily alter without flooding the model with particularized input. A judge's decision to explore beyond that subset should be rationalized within the preregistered research plan described above. The version of the given model used for interpretation should be the most current version made available to the public without charge. In doing so, we can improve the transparency of the research by allowing public access, and improve utility by limiting the research to high-quality LLMs. Perhaps in a perfect world, exact copies of the models would be delivered into public control for interpretive use so as to avoid tampering by the owners.<sup>117</sup> Though, of course, in that unlikely scenario, interpreters are deprived of future improvements to the models and must address the misalignment between the models available to the public and those employed by the courts.<sup>118</sup>

### 3. *Prompting Within the Same Conversational String*

In analyzing *Snell*, both of the prompts Judge Newsom put to the models were likely entered into the same chat. This likely resulted in

---

<sup>117</sup> It seems today any discussion of AI would be incomplete without recognizing a new player on the field, DeepSeek. DeepSeek is a Chinese AI model that is open-source, meaning anyone may download its code, and was reportedly trained at a fraction of the cost of other models. See Charlotte Edmond, *What is Open-Source AI and how Could DeepSeek Change the Industry?*, WORLD ECON. F. (Feb. 5, 2025), <https://www.weforum.org/stories/2025/02/open-source-ai-innovation-deepseek/> [<https://perma.cc/WXY7-3T9B>]. On one hand, an open-source model is very similar to the scenario in which the model is delivered into public control. On the other, if training requirements drop significantly, thereby making models easier to influence, the rationale for using a common selection of models based on the burden of their alteration weakens dramatically. In any case, using foreign-based software for interpretation could carry significant risks.

<sup>118</sup> AI companies are already developing and releasing tools for governmental use, seemingly targeted at providing exacting security features. See *Ross v. United States*, 331 A.3d 220, 231-32 (D.C. 2025). ChatGPT has even offered a government-focused product which functions on agencies' cloud servers. See J. Scott Schlegel, *The Wait May Be Over: Government AI Products Could Give Courts the Green Light*, [SCH] LEGAL TECH (Feb. 5, 2025), <https://judgeschlegel.substack.com/p/the-wait-may-be-over-government-ai> [<https://perma.cc/5KY5-7TD4>]. The efforts to protect sensitive data against exposure due to AI use is encouraging. However, these fixes seem solely directed at controlling what information goes out, rather than the quality of the information being accessed by the court. Moreover, the idea of a model being tailored to judicial use while still under private control may well undermine whatever safeguards exist by merit of broad accessibility to the public.

the second response being heavily impacted by the initial response due to the model's conversational memory. In essence, the model will utilize the information in its previous responses (as well as the user's previous prompts) to provide more relevant information as the conversation progresses.<sup>119</sup> Due to the nature of conversational memory in GenAI models, both ChatGPT and Bard likely placed a heavy emphasis on the ordinary meaning provided within the same string of prompts. Careful readers of ChatGPT's responses will notice parallels in the language of the responses to both prompts. Phrases like "altering the visible features of an [outdoor] area ... for aesthetic or practical purposes" and "[enhance/modify] the appearance and function[ality] of the [outdoor] space" are nearly identical.<sup>120</sup> Bard's similarities were less noticeable, but still present.<sup>121</sup> The ordering of the prompts has likely injected significant bias into the model's results by tacitly accepting the model's initial assertion. Perhaps Judge Newsom's clerk did access separate strings or clear the memory of the models between prompts, rendering those similarities coincidental. The matter is unclear. If Judge Newsom intended to apply the definition provided, prompting within the same chat would be a valid technique. If, however, Judge Newsom hoped for a "fresh" take on the question of whether installing an in-ground trampoline is landscaping, additional steps would have to be taken.

### ***Recommendation 3: Using Separate Chats; Clearing the Model's Memory***

The costs and benefits of this recommendation are intuitive. By routinely using separate conversational strings between prompts, one can prevent the adulteration of new data by previous responses, thus increasing the utility of that data in certain instances. Opposing practicality concerns include the additional time and resulting expense, and to a lesser extent expertise, required to manage this slightly more involved technique. One could easily overlook this step, prompt within the same string, and unknowingly access the model's conversational memory. This concern is mitigated somewhat by the other recommendations for preregistered research plans and exporting the results of the investigation (discussed within the analysis of *Deleon*). A word of caution: my experience with recent ChatGPT models has shown that the models can record information for access between strings as well. Manually clearing the model's memory would serve as an additional safeguard against this.

---

<sup>119</sup> For examples, see *infra* Appendix B.

<sup>120</sup> See *Snell*, 102 F.4th at 1234.

<sup>121</sup> *Id.* at 123435.

This isn't to suggest conversational memory is damaging without exception. There are instances in which the researcher would benefit from relying upon past responses (one of my following recommendations does just that). One instance in which prompting within the same string is acceptable is the application of context to a definition supplied by the model. Said context may include a request for a technical meaning as opposed to the ordinary meaning, imposing temporal or geographical limitations when generating definitions, and supplying the form of writing in which the term appears (insurance contract, statute, etc.).<sup>122</sup>

#### B. GENAI IN UNITED STATES V. DELEON

The facts of *Deleon* involve the defendant allegedly robbing a cashier at gunpoint.<sup>123</sup> Deleon was convicted of robbery and “brandishing a firearm during and in relation to a crime of violence,” two separate criminal counts.<sup>124</sup> The district court interpreted the sentencing guidelines to include Deleon’s actions as a form of “physical restraint” of the cashier.<sup>125</sup> Under those guidelines, Deleon’s sentence could be enhanced if his crimes involved an element of “physical restraint.”<sup>126</sup> The Eleventh Circuit Court of Appeals affirmed the district court’s finding.<sup>127</sup> Judge Newsom used this case as an opportunity to elaborate on his stance regarding GenAI for interpretation and to consider whether Deleon’s use of the gun constitutes “physical restraint.”<sup>128</sup>

Judge Newsom’s approach to using GenAI in *Deleon* is initially quite similar to his approach in *Snell*, though it quickly grows significantly more intricate. To begin, the Judge prompted ChatGPT (model GPT4-o) to provide “the ordinary meaning of ‘physically restrained.’”<sup>129</sup> Judge Newsom entered an identical prompt into Claude 3.5 Sonnet, an AI chatbot created by Anthropic,<sup>130</sup> twice within the same chat.<sup>131</sup> He went on to provide that same prompt—“What is the ordinary meaning of ‘physically restrained’?”—to three

---

<sup>122</sup> Basile, *supra* note 93, at 135, 139 (describing technical meaning as “non-ordinary or specialized” meaning).

<sup>123</sup> See *United States v. Deleon*, 116 F.4th 1260, 1274-76 (11th Cir. 2024).

<sup>124</sup> *Id.* at 1262.

<sup>125</sup> *Id.*

<sup>126</sup> *Id.*

<sup>127</sup> *Id.* at 1265.

<sup>128</sup> *Id.* at 1270-72.

<sup>129</sup> *Id.* at 1272.

<sup>130</sup> *Meet Claude*, ANTHROPIC, <https://www.anthropic.com/claude> [<https://perma.cc/ZMS8-EWFN>].

<sup>131</sup> *Deleon*, 116 F.4th at 1272-73.

separate GenAI models (ChatGPT, Claude, and Gemini) ten times each, with each prompt being in a new chat.<sup>132</sup> Additionally, one of Judge Newsom’s clerks provided the prompt to Open AI Playground, a paid model which allows for the direct adjustment of the model’s temperature.<sup>133</sup> The law clerk in this instance set the temperature to zero, theoretically minimizing the variability in Playground’s response.<sup>134</sup>

The analysis of Judge Newsom’s methodology in *Deleon* will follow a similar structure to the preceding analysis of *Snell*, in that each element discussed will be followed by recommendations for improvement. The following Sections begin by discussing recurring elements from *Snell*, including prompt clarity and conversational memory. They then analyze new elements of the interpretive strategy, including Judge Newsom’s attempt at large-batch prompting, and reducing the temperature of a given response. Each element is followed by recommendations for improvement. The recommendations include a novel large-batch prompting process, the generation of a response with minimized temperature using a consistent model, and the exportation of results for ex-post verification.

1. *Recurring Topics from Snell: Prompt Clarity; Conversational Memory*

Once again, the clarity of the prompt is worth noting. “What is the ordinary meaning of ‘X’?” will do little to promote or discourage variety in the responses. This prompt is also free of any surplus information which might guide the model to select any one given meaning. For the sake of consistency across users and as an introduction to the research, this is a creditable beginning.

Judge Newsom’s methodology offers a new wrinkle: a model was provided an identical prompt within the same string resulting in similar responses that differed in length. The repetition raises some interesting considerations. As discussed previously, and as Judge Newsom notes, by prompting the model a second time within the same chat, the Judge likely employed Claude’s conversational memory.<sup>135</sup> We then face a problem reminiscent of the application of the ordinary meaning provided in *Snell*. Specifically, the model potentially “doubled down” on previous assertions. The model’s

---

<sup>132</sup> *Id.* at 1273–74.

<sup>133</sup> *Id.* at 1274 n.5.

<sup>134</sup> *Id.*

<sup>135</sup> *Id.* at 1276 n.10 (referring to the topic addressed in this Note as “conversational history”).

probabilistic generation may still produce ideas not originally captured by the initial response, but the new output will almost certainly be partially reliant on the context of the conversation. It is unclear to me why the second response would be markedly shorter than the initial one.<sup>136</sup> It may be the model is programmed to tend toward the succinct in response to multiple identical queries in order to preserve computational resources for other users. Or it may be that multiple passes at the same question with the context of a prior response results in a more refined, and therefore shorter, output. I think the similarity in specific word combinations we see between the two results could indicate either, or a combination of the two.<sup>137</sup> It may even be that the repeated prompts indicate to the model that the user didn't understand the previous response, encouraging a more simplistic answer. In any case, those using repetitive prompts for research pursuits should be conscious of using distinct prompt strings and clearing the model's memory between prompts to avoid this informational impurity.

## 2. Judge Newsom's Large-Batch Prompting Attempt

The most significant deviation from the approach in *Snell* is the use of 10 repetitive promptings, each within a new conversational string, on three separate models. The primary benefit of utilizing separate conversational strings is that it allowed Judge Newsom to avoid the conversational memory pitfall we see above. By separating the prompts, Judge Newsom was thereby able to access a wider array of potential responses, all fairly likely according to the probabilities determined by the model. The higher number of prompts serves several purposes. First, a variety of responses mitigates the possibility of an outlier response being selected and presented as the model's best attempt. Second, it affords the user an opportunity to identify consistency between responses, which might constitute a conceptual "common core."<sup>138</sup> That "common core," if we accept Judge Newsom's understanding, may serve as evidence of ordinary meaning.<sup>139</sup>

The third and, in my perspective, most consequential benefit of this large-batch prompting method is its use in revealing aspects of ordinary meaning which may be reasonably common, yet far from ubiquitous. For example, Judge Newsom's interactions with Claude reveal a theme of categorizing physical objects being introduced to

---

<sup>136</sup> *Id.* at 1272–73.

<sup>137</sup> *Deleon*, 116 F.4th at 1272–73.

<sup>138</sup> *Id.* at 1274–77.

<sup>139</sup> *Id.* at 1276.

“block one’s path” as a form of “physical restraint.”<sup>140</sup> Precisely these words are used in four of the model’s responses (Responses 1, 4, 6, and 9). Some other responses include sufficiently general language to encompass this more specific example. However, as those who study judicial interpretation can attest, specificity carries with it an interpretive strength that the general does not.<sup>141</sup>

It is the relevance of these “common but not too-common” specifics to ordinary meaning on which I suspect Judge Newsom and I may disagree. One with an eye solely toward the “common core” of the various responses might naturally cleave away this repeated, but hardly essential, notion as creative “fuzz,” which distracts from, rather than illuminates, the ordinary meaning.<sup>142</sup> The better response, I think, is to recognize that even those ancillary patterns that arise from large-batch prompting may have significant bearing on ordinary meaning.<sup>143</sup> “Blocking one’s path” may not be what most consider to be physical restraint, but its explicit recurrence in the model’s responses may serve as evidence many do. A judge may prudently determine these results require further and more specific evidence, such as legislative history, to resolve this apparent lack of clarity.

I want to stress the importance of this approach and how it differs from Judge Newsom’s efforts in *Snell*, and even *Deleon*. By moving to larger numbers of prompts and by looking for connections between them, we have moved beyond attempting to use GenAI as a sort of changeable dictionary. I am not asserting that the identical content between the responses *is* the ordinary meaning, as Judge Newsom suggests it might be.<sup>144</sup> Instead, GenAI is being used to flag essential elements of ordinary meaning which might escape the intuitions of the interpreter. In doing so, the variability of GenAI

---

<sup>140</sup> *Id.* at 1279–82.

<sup>141</sup> VALERIE C. BRANNON, CONG. RSCH. SERV., R45153, STATUTORY INTERPRETATION: THEORIES, TOOLS, AND TRENDS 55 (2023) (providing a description of the general/specific canon).

<sup>142</sup> *See Deleon*, 116 F.4th at 1277 (explaining that Judge Newsom finds it sensible that LLMs would generate outputs that are “a little... fuzzy around the edges” given their training on “uses of language in the real world”).

<sup>143</sup> This stance remains predicated on the threshold assumption that GenAI’s sweeping collection, analysis, and presentation of data remains pertinent to ordinary meaning despite the practical issues described in Part IV.

<sup>144</sup> *See Deleon*, 116 F.4th at 1276.

responses transforms from a liability into an asset, as we seek out the imprecise contours of ordinary meaning.

It is worth mentioning a countervailing principle. Some may argue it is inappropriate for the judge to consider GenAI-informed aspects of the interpretive question *sua sponte* as a matter of respecting the adversarial process. This principle might caution the judge to ignore those patterns first observed by the interpreter. That said, to ignore those patterns does seem to intentionally weaken the interpretation itself. Moreover, this issue seems analogous to the selection between dictionaries as a function of inadequately captured definitional elements.<sup>145</sup>

#### ***Recommendation 4: Large-Batch Prompting Investigative Process***

The purpose of this large-batch prompting process is to begin pinning down Judge Newsom’s “common core” of meaning, as well as other relevant patterns, among an assortment of responses. This is done by quickly generating a wide variety of plausible ordinary meanings (Step 1), sorting them to identify commonalities between them by applying simple factual context (Step 2), and applying a textual context to look for further patterns revealed by shifts in their categorization (Step 3). The method I suggest in this Section is merely one among many potentially beneficial research processes. In fact, it may be more appropriate to conceive of this process as an example of a preregistered research plan, or a general appeal to standardization, as opposed to a recommendation all its own. However, I believe this three-step process reasonably balances the interests in generating a large sample size to account for and utilize the variation in GenAI responses, respecting the judge’s limited time and resources, and considering basic facts and textual context. This process is meant to be quick and simple to employ. Large-batch prompting can be done by any individual with access to current, free versions of high-quality GenAI models. Researchers would require only a basic understanding of the models themselves. The process would be fairly standardizable across cases. I believe it has the potential to shed light on relevant inter-definitional patterns for the interpretive process. In summary, it relates favorably to the transparency, utility, and practicality concerns discussed above. In Appendix B, I have provided two examples of this process as applied to the circumstances of *Snell* and *Deleon*.

---

<sup>145</sup> See Scalia & Garner, *supra* note 68, at 420–21 (urging the use of some dictionaries which consider an “all-important element” of a definition, where others do not).

Nevertheless, this process is subject to substantial limitations. It fails to consider a phrase's use within a given sentence, as part of a list, or otherwise. No portion of the process directly addresses when a technical meaning, as opposed to an ordinary meaning, is called for, though it may indirectly address the issue within the third step. The process also does not consider other external evidence of meaning, such as legislative history or custom. The nature of the prompt construction also allows for a small measure of discretion in describing some basic facts within the prompt, thereby introducing the opportunity for transparency-related concerns. An issue not unique to this process, but inescapably important, is the potential for models to fail to adequately follow directions. Indeed, in performing my *Snell* analysis, ChatGPT failed to account for one definition (Definition 16) under Step 3. In the spirit of following my own "research plan," these original, incomplete results were used in favor of generating a new response (See Appendix B.1). One absent analysis does not substantially detract from the pattern-based value of the method, but it does highlight the potential for incomplete responses. If ChatGPT's response features a particularly egregious omission, it may be appropriate to prompt for a response to the excluded definitions directly. For example, "Perform this analysis for Definitions 2, 4, and 6." Another example of ChatGPT deviating from instructions can be seen in Appendix C. The model consistently summarized (rather than quoted) its own definitions. While not problematic if the original definition were still applied to the Step 2 analysis, for Definitions 3 and 7, the model clearly applies its own summary, as opposed to the original definition, to perform the analysis. Step 2 and 3 responses featuring summarized or paraphrased definitions should thus be thoroughly scrutinized to determine whether the underlying logic actually relates to the definitions supplied. All this to say, this process, in isolation, is insufficient for determining interpretive questions.

This process is meant to be engaged in by the interpreter, and not necessarily other parties. Perhaps a judge would be interested in having opposing parties present common core analyses. In that case, the judge should be wary of the absence of preregistration or other safeguards and should be prepared to find meaningful distinctions between the proposals, likely in the frequency and specificity of the patterns identified.

*Step 1: Generating Ordinary Meanings*

Prompt 1: What is the ordinary meaning of “X”? Please respond twenty times, at a length of three sentences each.

The first step of this process shares some similarities with Judge Newsom’s approach in *Deleon*. It involves asking for multiple iterations of responses to the question, “What is the ordinary meaning of X?” However, instead of providing the prompt in ten separate conversational strings, the user will generate each of the definitions within a single prompt. In an effort to avoid the shortening of the responses we see with Judge Newsom’s repeated prompts to Claude, the length of each definition will also be provided.<sup>146</sup> Additionally, the prompt calls for twenty responses as opposed to ten, in the hopes of garnering a wider variety of responses and, theoretically, improving utility by alerting the user to more dimensions of the interpretive question.

Undoubtedly, given my previous discussion of conversational memory, there will be some confusion regarding my decision to call for the definitions within the same conversational string. I came to this decision for two reasons. The first relates to practicality. Accessing a new string for each prompt quickly becomes cumbersome. As the number of responses increases, the time dedicated to the task will accumulate quickly for the busy judge. Secondly, my limited research with the method in no way indicated that the later definitions were dramatically influenced by those earlier in the response, as the user has had no chance to engage with them. Moreover, having multiple responses within the same prompt might unintentionally encourage our now valuable variation, rather than diminish it, as the model seeks to provide novel phrasings of the definition. This remains theoretical, and may, of course, be proven wrong.

*Step 2: Application of Simple Facts*

Prompt 2: Please list all definitions, of those 20 you listed above, under which you would consider (Simple Facts Y: Ex. constructing a below-ground trampoline / a cashier held at gunpoint) to be “X.” Explain why. For those you would not consider to be “X,” explain why.

---

<sup>146</sup> See *Deleon*, 116 F.4th at 1272–73.

The second step of this process is an effort to determine patterns through categorization. In essence, we force the model to include or exclude the term for each definition and demand a rationale. For close questions of judicial interpretation, it is unlikely that all 20 of the definitions are going to point strongly in favor of inclusion or exclusion. Frankly, even if the results were to uniformly suggest a given result, the interpreter should meet the analysis with the same level of scrutiny. It bears mentioning that the framing of the “simple facts” offers a measure of discretion. Interpreters should endeavor to keep the facts brief and avoid subjective phrasings when providing the context as a matter of transparency.

The intent of this step is not to provide the interpreter a series of rationales from which to choose. Instead, it allows the interpreter to quickly identify patterns relevant to the interpretive question without requiring an assessment of 20 definitions. Taking my application of the second step to *Snell* as an example (See Appendix B.1), even a cursory reading of the response to my second prompt will reveal commonly recurring considerations related to the ordinary meanings of landscaping. These include improving functionality as opposed to solely aesthetic appeal, and changing the natural versus man-made nature of the land’s elements. An industrious interpreter may be forced to return to the Step 1 results to determine less obvious, but potentially relevant, patterns. The express inclusion of water features (See Appendix B.1 Definitions 2,7,8,10,13,14), were they germane to the trampoline issue, serves as an example.

With those patterns identified, the interpreter is left to ponder the extent to which those patterns inform the ordinary meaning. There is an unavoidable element of subjectivity in that determination. The question of the precise bearing any given pattern has on the determination of ordinary meaning exceeds the scope of this Note. Suffice it to say, specific and common patterns may be stronger evidence for inclusion/exclusion than general and infrequent patterns. I would caution any potential interpreter against the assumption that because the GenAI model considered 13 of the 20 definitions generated for “landscaping” to include a below-ground trampoline, those interpretations should be considered strong evidence for inclusion. Were the model prompted exactly this way a second time, the analytical conclusions would likely differ. I have included an example of a second set of responses for Step 2 within Appendix C. This method relies on the perceived strength of GenAI to detect patterns, not to perform thoughtful analysis.

*Step 3: Application of Broad Textual Context*

Prompt 3: Assume “X” in this context was used within a[n] (Textual Context: Ex. insurance policy / criminal statute). Would your answer as to whether Y is “X” change for each of the 20 definitions you have provided?

This final step serves as a rudimentary attempt to elicit further patterns related to ordinary meaning by providing some level of context. I have selected this form of textual context because the question can be standardized for virtually any written work. Transparency of the process is thus improved by reducing the subjectivity involved.

Once again, the purpose of this step is not for the interpreter to select between given rationales. Nor is the purpose to provide any level of statistical analysis based on the resulting categorization under the ordinary meanings provided. The purpose is merely to shed insight on potentially relevant conceptual patterns which might arise from the introduction of this base textual context.

Looking at my application of this process to the facts of *Snell*, the Step 3 results reveal a subtle potential pattern, that of “visibility.” The model raises the issue of whether “landscaping” is a concept that is tied to alterations designed to be visible.<sup>147</sup> A below-ground trampoline, while certainly a functional alteration, is likely designed to avoid the eyeline of a casual onlooker. This factor is hardly dispositive, but it may be relevant to the interpretive matter. Without the introduction of this textual context, this peripheral facet of the interpretive question could have escaped the attention of the interpreter.

It should be obvious that this single prompt does not approach a complete introduction of the relevant context for any interpretive question. The prompt omits factors like whether the term is used as part of a list, the title of the section in which the term is listed, other definitions provided within the text, and historical context. I remain open to the idea of introducing further context for pattern identification. Doing so could yield interesting insights. However, it would also increase the subjectivity of the research process by muddying the standardization of the prompting method.

---

<sup>147</sup> See *infra* Appendix B.1.

### 3. *Response with Minimized Temperature*

A final praiseworthy (and understated) element of Judge Newsom’s approach in *Deleon* is his law clerk’s work generating a response with minimized temperature.<sup>148</sup> A generous interpretation of this response as the model’s probabilistic “best guess” could recommend it as a form of weak evidence of ordinary meaning in isolation, given the more deterministic results.<sup>149</sup> However, this seems to be another attempt to utilize a changeable source as a form of dictionary, a use that is at odds with this Note’s understanding of GenAI’s potential interpretive benefits.

Additionally, however, a response with minimized temperature could serve as a sort of “control condition” when compared to large-batch prompting. The user might consider patterns which vary from the control with increased skepticism, and more readily acknowledge large-batch results that are aligned. Though, for reasons discussed above, I would caution interpreters against discounting potentially significant patterns on this basis alone.

The issue with the application of this technique in the context of *Deleon* is the variation between models. OpenAI Playground might provide a reasonable control for prompts entered into ChatGPT. The models were created by the same company, and were likely trained using similar corpora.<sup>150</sup> The low-temperature response provided by OpenAI Playground may have no bearing on responses from Claude and Bard, as these LLMs were likely trained differently.

#### ***Recommendation 5: Response with Minimized Temperature Using a Consistent Model***

Assuming a response with minimized temperature does offer utility benefits as a control condition, researchers may wish to provide a response with minimized temperature for each model they choose to access. Unfortunately, free models typically do not allow users to modify temperature directly. Given that part of the transparency value of GenAI is the accessibility of the technology, other steps might have to be taken to provide similar results without imposing a financial cost. To that end, those who wish to approximate a minimized temperature response without access to

---

<sup>148</sup> See *Deleon*, 116 F.4th at 1274 n.5.

<sup>149</sup> *Id.*

<sup>150</sup> Mike Paul, *Exploring the Limitless Possibilities: OpenAI Playground vs. ChatGPT*, TECHPILOT.AI (March 15, 2024), <https://techpilot.ai/openai-playground-chat-gpt-playground/> [<https://perma.cc/8GKS-DCTF>].

paid versions may wish to introduce alterations to their prompts. Examples of these alterations could include “Please respond at the lowest possible temperature level,” “Prioritize certainty over creativity,” or “Provide your best guess.” Having experimented with this technique, it seems to modestly reduce variability in the responses. I have included my (admittedly limited) results in Appendix D.

This recommendation does call for the alteration of an individual parameter. As a matter of practicality, it is therefore more difficult to perform than many of the other recommendations. This is particularly true in light of the need to substitute prompts for direct parameter alteration when accessing free models. For these reasons, there is certainly room for argument that the utility benefits of this approach are outweighed by the additional complexity.

#### ***Recommendation 6: Exporting Results for Ex-Post Verification***

Mandating the exportation of results serves a basic transparency-related function. Attorneys and members of the public would be able to see, and theoretically approximate, responses generated during the interpretive process. Judge Newsom provided a reasonable facsimile of this concept by including detailed notes of his investigation. However, exporting the data to be stored as a digital file, thereby providing a clear transcript of the research process, allows interested parties to quickly compare the results with the preregistered research plan and detect deviation. In considering the practicality of this suggestion, executing this technique will require a higher level of experience with the given model. Still, exporting the results can be done fairly easily.

The exported results may still be subject to manipulation by sophisticated parties. However, no research method is entirely immune to meddling, and I find it relatively unlikely that a judge willing to use GenAI, develop a preregistered research plan, and subsequently export the results would then engage in technical tampering to fit their intuitions. Exporting results is meant to serve as an additional layer of reassurance for the public and to reinforce the credibility of judicial results, not to ensure the results are unquestionably accurate. Were judges to allow parties to submit their own GenAI results, requiring those parties to export and submit complete, consolidated accounts of their inquiries seems a reasonable demand.

## VI. CONCLUSION

Judge Newsom's inaugural efforts to apply artificial intelligence to interpretation in *Snell* and *Deleon* prove both fascinating and challenging. Is the common core between responses the ordinary meaning, as Judge Newsom suggests?<sup>151</sup> What problems and benefits derive from this use of a revolutionary technology? How does GenAI comport with the reigning interpretive scheme?

In my opinion, Judge Newsom's analysis as to why GenAI predictive responses (as patterned off of extensive and varied use of language) could inform ordinary meaning makes intuitive sense and, in my eyes, matches the value any contemporary dictionary could provide. However, the stark reality of the practical issues with GenAI today, including hallucinations and the private control over the models to be accessed, should preclude use of GenAI for interpretation on a widespread basis. While these practical concerns can be mitigated to varying extents, taken together, they present too significant a danger of unreliability to justify adoption. If, despite my misgivings, GenAI is so used, it should be used in conjunction with rigorous transparency standards and guidelines on research strategy. To that end, my recommendations include mandated preregistered research plans, LLM standardization, using separate conversational strings and clearing memory where appropriate, taking advantage of the benefits of large-batch prompting, generating low-temperature responses on consistent models, and the exportation of results mentioned above. Moreover, I believe the simple application of a generated definition is insufficient to shed light on the interpretive question under even the best circumstances. Instead, GenAI will have to be marshalled more creatively, so as to align its strengths with those of the interpreters. One potential application is the flagging of conceptual patterns related to the issue and allowing the interpreters to weigh the value of those patterns. "Blocking one's path" as physical restraint in *Deleon* serves as an example. Effective use of GenAI for interpretation will require not just a set of careful boundaries on how results are generated, but also a healthy dose of the reality that GenAI should not be substituted for common sense and extensive research.

For those who subscribe to textualist thinking for purposes of interpretation, this Note serves as an introduction on how GenAI may serve the espoused motivations of textualist ideology. While a far cry from the perfect interpretive tool, GenAI offers some tangible benefits over existing methods. It is readily accessible to the public,

---

<sup>151</sup> See *Deleon*, 116 F.4th at 1276.

reflective of the common understanding of meaning (at least theoretically), and capable of dealing with context to a degree unavailable with dictionaries. If GenAI truly can assist interpreters in making more decisions in line with what drafters effectuated, then perhaps we will see more predictable judicial results, fair notice, and legal stability.

This topic is a challenging one. Undoubtedly, I have overlooked some important issues due to my own limited understanding of GenAI, textualism, and legal interpretation more broadly. My desire in writing this Note was not to criticize Judge Newsom's approach in either *Snell* or *Deleon*. My hope was to build on his approach and explore its other implications. I am grateful that Judge Newsom saw fit to address the issue twice, with obvious consideration and research in each instance.

APPENDICES<sup>152</sup>**Appendix A: “Creative” Results****User**

What is the ordinary meaning of “landscaping”? Respond five times at a length of three sentences each. Be very creative in your responses.

**ChatGPT**

1	Landscaping is the art of shaping and transforming outdoor spaces to create visually appealing environments. It often involves planting trees, designing pathways, and arranging natural elements like rocks and water features. In its essence, it’s the fusion of architecture and nature, turning plain land into something that sparks joy and harmony.
2	Landscaping refers to the deliberate design and cultivation of outdoor areas to improve aesthetic appeal. This can include everything from planting lush gardens to installing patios and sculptures that blend seamlessly with nature. It’s a creative way to balance beauty and functionality, inviting people to experience the outdoors in an intimate way.
3	In its simplest form, landscaping is the act of modifying the terrain to make it more attractive and functional. This may include leveling ground, adding plant beds, or incorporating unique features like fountains or pergolas. It’s about crafting a space that not only looks beautiful but also provides a peaceful retreat.
4	Landscaping involves shaping the natural environment to align with human preferences, blending nature’s raw elements with thoughtful design. It’s more than just planting flowers or trees—it’s about creating a cohesive outdoor narrative that flows from one element to the next. Whether it’s a serene garden or a vibrant urban park, landscaping breathes life into spaces.

---

<sup>152</sup> The responses featured in the appendices have been edited for formatting. The content, however, remains true to the information supplied by the LLM. The summaries and elaborations included at the end of some responses were authored by the LLM as part of those responses.

<b>5</b>	Landscaping can be described as nature’s makeover, where soil, stone, and plants are rearranged into patterns of beauty. It’s like painting a canvas with greenery and hardscapes, using the land itself as the medium. A well-landscaped area can evoke emotions, inspire creativity, and provide sanctuary from the hustle and bustle of life.
----------	--

## Appendix B: Large-Batch Prompting for *Snell and Deleon*

### Appendix B.1: *Snell*<sup>153</sup>

#### User

What is the ordinary meaning of “landscaping”? Please respond twenty times, at a length of three sentences each.

#### ChatGPT

<b>1</b>	Landscaping refers to the process of modifying and designing the visible features of a piece of land, especially around buildings and gardens. It typically includes planting trees, shrubs, and flowers, as well as installing pathways, ponds, and other decorative elements. The goal is to create a visually appealing and functional outdoor space.
<b>2</b>	Landscaping involves the art and science of planning, designing, and maintaining land areas for aesthetic or functional purposes. It often includes elements like lawns, plants, hardscaping (like patios or retaining walls), and water features. This practice enhances the beauty and usability of a space, whether residential or commercial.
<b>3</b>	Landscaping is the activity of altering a piece of land to make it more attractive or practical. This could involve planting plants, building structures like decks or walls, and arranging outdoor features. It’s often done to improve the environment around homes, businesses, or public spaces.
<b>4</b>	Landscaping focuses on creating and maintaining an appealing outdoor environment by modifying natural elements such as plants and terrain. It also involves incorporating man-made

---

<sup>153</sup> As mentioned above, ChatGPT actually omitted an analysis for Definition 16 within the Step 3 results. *See supra* p. 576. This should serve as a cautionary sign to avoid total reliance on analytical or statistical determinations by the model.

	features like walkways or lighting. Landscaping serves both decorative and functional purposes, often enhancing the property's overall value.
<b>5</b>	Landscaping refers to the process of designing and transforming an outdoor space to improve its appearance and functionality. It can involve both natural elements, like plants and trees, and constructed features, like patios or fountains. The primary aim is to create an inviting and harmonious environment.
<b>6</b>	Landscaping is the practice of modifying land to make it more aesthetically pleasing, often by arranging plants, trees, rocks, or other natural elements. This can also involve creating paths, seating areas, or installing lighting to enhance the space. The purpose is usually to create a serene or visually striking environment.
<b>7</b>	Landscaping involves shaping the land by adding or rearranging elements like trees, flowers, rocks, or water features to improve the overall look. It also includes hardscaping features like fences, patios, or gazebos. This practice is common in both residential and public areas to enhance outdoor aesthetics.
<b>8</b>	Landscaping is the design and organization of outdoor areas, incorporating both living plants and constructed elements like paths, walls, or water features. This process aims to beautify and make outdoor spaces more functional. It often enhances the natural beauty of the environment, adding value to the property.
<b>9</b>	Landscaping refers to the intentional arrangement and modification of outdoor spaces for both beauty and function. It commonly involves planting trees, flowers, and shrubs while adding features like walkways, seating areas, and fountains. The goal is to create a pleasing and sustainable environment.
<b>10</b>	Landscaping is the process of designing and maintaining an outdoor area by altering its natural features, such as planting new vegetation or adding stone pathways. It may also involve building retaining walls or water features. This art form enhances both the appearance and utility of the space.
<b>11</b>	Landscaping involves enhancing the outdoor environment by designing and implementing plantings, ground cover, and

	<p>other natural features. In many cases, it also includes the construction of hardscapes, such as patios, walls, or decks. Landscaping makes outdoor areas more pleasant and usable for people.</p>
<b>12</b>	<p>Landscaping is the design and organization of outdoor areas to create a visually appealing and functional environment. It often includes the planting of flowers, shrubs, and trees, as well as the construction of hardscape elements like patios or paths. Landscaping enhances both the aesthetic and practical aspects of a property.</p>
<b>13</b>	<p>Landscaping refers to the practice of arranging outdoor spaces by adding natural and man-made elements to improve their visual appeal. This may include planting gardens, installing water features, or building terraces. The main goal is to create an attractive and comfortable environment for people to enjoy.</p>
<b>14</b>	<p>Landscaping is the art of transforming outdoor spaces by adding plants, rocks, water features, and other elements to improve aesthetics and functionality. It may involve creating new pathways, constructing retaining walls, or installing lighting. The aim is often to make the outdoor area both practical and visually striking.</p>
<b>15</b>	<p>Landscaping is the process of planning and organizing outdoor spaces to enhance their natural beauty and usability. It typically involves planting trees, shrubs, and flowers, as well as adding features like ponds, patios, or walkways. Landscaping is designed to create attractive, enjoyable outdoor environments.</p>
<b>16</b>	<p>Landscaping involves the alteration and maintenance of outdoor areas through the strategic arrangement of plants, rocks, and other decorative elements. It can include both softscaping, like flower beds, and hardscaping, like stone walls or pathways. The goal is to enhance the property's overall aesthetic and functionality.</p>
<b>17</b>	<p>Landscaping refers to the planning and design of outdoor spaces by incorporating both natural and constructed features. This can include planting trees, building paths, and installing water elements to create an inviting environment. Landscaping is often used to boost the appeal of homes, parks, or commercial properties.</p>

<b>18</b>	Landscaping is the practice of improving the appearance and usability of outdoor spaces through the strategic placement of plants, trees, and constructed features. It often involves creating a balance between natural elements and hardscaping, such as fences or patios. The result is a more harmonious and attractive outdoor area.
<b>19</b>	Landscaping focuses on enhancing the visual appeal and functionality of outdoor environments through the placement of plants, rocks, and man-made structures. It can involve projects like installing gardens, constructing terraces, or adding decorative lighting. The purpose is to create outdoor spaces that are both beautiful and practical.
<b>20</b>	Landscaping refers to the design and modification of land areas to improve their aesthetic or practical use. This can involve planting greenery, arranging stones or other decorative materials, and installing features like water fountains. Landscaping is often aimed at creating a peaceful, well-organized outdoor environment.

### User

Please list all definitions, of those 20 you listed above, under which you would consider constructing a below-ground trampoline to be “landscaping.” Explain why. For those you would not consider to be “landscaping,” explain why not.

### ChatGPT

Here’s how I would categorize the definitions regarding whether constructing a below-ground trampoline would be considered “landscaping”:

Definitions that could include constructing a below-ground trampoline as “landscaping”:

<b>1</b>	<p>Definition 3:</p> <p>“Landscaping is the activity of altering a piece of land to make it more attractive or practical.”</p> <p>- A below-ground trampoline could be considered practical and an aesthetic improvement, aligning with the idea of making the space more functional and attractive.</p>
----------	--

<b>2</b>	Definition 5:  “Landscaping refers to the process of designing and transforming an outdoor space to improve its appearance and functionality.”  - Installing a below-ground trampoline transforms the space and enhances functionality, making it a possible form of landscaping.
<b>3</b>	Definition 9:  “Landscaping refers to the intentional arrangement and modification of outdoor spaces for both beauty and function.”  - The trampoline adds functionality and could enhance the beauty of the space if integrated smoothly into the design.
<b>4</b>	Definition 10:  “Landscaping is the process of designing and maintaining an outdoor area by altering its natural features, such as planting new vegetation or adding stone pathways.”  - Although this mentions pathways and vegetation, “altering its natural features” could encompass creating a trampoline pit.
<b>5</b>	Definition 11:  “Landscaping involves enhancing the outdoor environment by designing and implementing plantings, ground cover, and other natural features.”  - While focused on plants and ground cover, the addition of a trampoline could be seen as part of designing the outdoor environment.
<b>6</b>	Definition 12:  “Landscaping is the design and organization of outdoor areas to create a visually appealing and functional environment.”  - A trampoline could contribute to both function and visual appeal, fitting this definition.
<b>7</b>	Definition 13:  “Landscaping refers to the practice of arranging outdoor spaces by adding natural and man-made elements to improve their visual appeal.”

	- A trampoline is a man-made element that could enhance both visual appeal and usability.
<b>8</b>	<p>Definition 15:</p> <p>“Landscaping is the process of planning and organizing outdoor spaces to enhance their natural beauty and usability.”</p> <p>- Usability is key here, and a trampoline definitely adds functionality to the space.</p>
<b>9</b>	<p>Definition 16:</p> <p>“Landscaping involves the alteration and maintenance of outdoor areas through the strategic arrangement of plants, rocks, and other decorative elements.”</p> <p>- The trampoline could be considered a decorative or functional element in this case.</p>
<b>10</b>	<p>Definition 17:</p> <p>“Landscaping is the planning and design of outdoor spaces by incorporating both natural and constructed features.”</p> <p>- A trampoline is a constructed feature, and its incorporation could fit into the design of a landscape.</p>
<b>11</b>	<p>Definition 18:</p> <p>“Landscaping is the practice of improving the appearance and usability of outdoor spaces through the strategic placement of plants, trees, and constructed features.”</p> <p>- The trampoline is a constructed feature that improves usability, aligning with this definition.</p>
<b>12</b>	<p>Definition 19:</p> <p>“Landscaping focuses on enhancing the visual appeal and functionality of outdoor environments through the placement of plants, rocks, and man-made structures.”</p> <p>- A below-ground trampoline is a man-made structure that enhances functionality.</p>
<b>13</b>	<p>Definition 20:</p> <p>“Landscaping refers to the design and modification of land areas to improve their aesthetic or practical use.”</p>

	- The trampoline is practical and can contribute to aesthetics when designed well.
--	--

Definitions that would not include a below-ground trampoline as “landscaping”:

<b>1</b>	<p>Definition 1:</p> <p>“Landscaping refers to the process of modifying and designing the visible features of a piece of land, especially around buildings and gardens.”</p> <p>- Since this emphasizes “visible features,” a below-ground trampoline, which is recessed and not as prominent, may not qualify.</p>
<b>2</b>	<p>Definition 2:</p> <p>“Landscaping involves the art and science of planning, designing, and maintaining land areas for aesthetic or functional purposes.”</p> <p>- Although functional, this definition focuses more on traditional elements like plants and hardscapes, not something as unconventional as a trampoline.</p>
<b>3</b>	<p>Definition 4:</p> <p>“Landscaping focuses on creating and maintaining an appealing outdoor environment by modifying natural elements such as plants and terrain.”</p> <p>- This is more about natural elements, and a trampoline is not a natural modification of the terrain.</p>
<b>4</b>	<p>Definition 6:</p> <p>“Landscaping is the practice of modifying land to make it more aesthetically pleasing, often by arranging plants, trees, rocks, or other natural elements.”</p> <p>- It specifically mentions natural elements, and the trampoline, as a man-made object, doesn’t fit here.</p>
<b>5</b>	<p>Definition 7:</p> <p>“Landscaping involves shaping the land by adding or rearranging elements like trees, flowers, rocks, or water features to improve the overall look.”</p>

	- This definition is more focused on improving visual appeal with plants and water features, rather than something like a trampoline.
<b>6</b>	<p>Definition 8:</p> <p>“Landscaping is the design and organization of outdoor areas, incorporating both living plants and constructed elements like paths, walls, or water features.”</p> <p>- Though it mentions constructed elements, this focuses on traditional hardscaping (paths, walls), which a trampoline does not fit into.</p>
<b>7</b>	<p>Definition 14:</p> <p>“Landscaping is the art of transforming outdoor spaces by adding plants, rocks, water features, and other elements to improve aesthetics and functionality.”</p> <p>- This emphasizes plants, rocks, and water features, and while a trampoline adds functionality, it might not be seen as an improvement to aesthetics in the same way.</p>

### Summary:

13 out of 20 definitions could consider the construction of a below-ground trampoline as landscaping because they focus on improving usability, functionality, and often incorporate man-made or constructed features as part of landscaping.

The remaining 7 definitions would not classify it as landscaping, as they are more focused on natural elements like plants, trees, or traditional hardscaping like walls or pathways. These definitions emphasize modifying visible, natural features rather than something like a trampoline.

### User

Assume “landscaping” in this context was used within an insurance policy. Would your answer as to whether constructing a below-ground trampoline is “landscaping” change for each of the 20 definitions you have provided?

### ChatGPT

Yes, the context of an insurance policy could affect the interpretation of “landscaping,” as legal and contractual terms are

often interpreted more strictly or in line with common usage. Here's how my assessment would change for each of the 20 definitions, considering that insurance contracts usually require clarity and adherence to common understanding:

Definitions that still could include a below-ground trampoline as "landscaping" in an insurance policy:

<b>1</b>	<p>Definition 3:</p> <p>"Landscaping is the activity of altering a piece of land to make it more attractive or practical."</p> <p>- In an insurance context, practicality could still apply to a trampoline if it enhances the usability of the space. This definition is broad enough to include a functional item like a trampoline.</p>
<b>2</b>	<p>Definition 5:</p> <p>"Landscaping refers to the process of designing and transforming an outdoor space to improve its appearance and functionality."</p> <p>- The trampoline could qualify under "functionality," especially if the insurance policy emphasizes the utility of landscaping as part of outdoor space improvement.</p>
<b>3</b>	<p>Definition 9:</p> <p>"Landscaping refers to the intentional arrangement and modification of outdoor spaces for both beauty and function."</p> <p>- Since insurance contracts may focus on the function of the landscape, a trampoline could fall under this broad definition.</p>
<b>4</b>	<p>Definition 10:</p> <p>"Landscaping is the process of designing and maintaining an outdoor area by altering its natural features."</p> <p>- The focus on altering natural features can still apply to creating a trampoline pit. This might be acceptable, depending on how an insurer interprets "maintaining" outdoor areas.</p>
<b>5</b>	<p>Definition 12:</p> <p>"Landscaping is the design and organization of outdoor areas to create a visually appealing and functional environment."</p>

	<p>- If the insurance policy allows for a focus on functionality, the below-ground trampoline could be seen as an element of landscaping.</p>
<b>6</b>	<p>Definition 13:</p> <p>“Landscaping refers to the practice of arranging outdoor spaces by adding natural and man-made elements to improve their visual appeal.”</p> <p>- Since a trampoline is a man-made element, and the policy could interpret “man-made elements” broadly, this definition may still include a trampoline.</p>
<b>7</b>	<p>Definition 15:</p> <p>“Landscaping is the process of planning and organizing outdoor spaces to enhance their natural beauty and usability.”</p> <p>- Usability is still key, and if a below-ground trampoline is seen as improving usability, it might fit the definition under an insurance policy.</p>
<b>8</b>	<p>Definition 17:</p> <p>“Landscaping is the planning and design of outdoor spaces by incorporating both natural and constructed features.”</p> <p>- Since the trampoline is a constructed feature, it could still qualify in an insurance context, depending on how strictly “constructed features” is defined.</p>
<b>9</b>	<p>Definition 18:</p> <p>“Landscaping is the practice of improving the appearance and usability of outdoor spaces through the strategic placement of plants, trees, and constructed features.”</p> <p>- In an insurance policy, the trampoline, as a “constructed feature,” would likely still be considered landscaping if usability is the focus.</p>
<b>10</b>	<p>Definition 19:</p> <p>“Landscaping focuses on enhancing the visual appeal and functionality of outdoor environments through the placement of plants, rocks, and man-made structures.”</p>

	- The trampoline is a man-made structure, and its functionality could fit the scope of an insurance policy, depending on how broadly they interpret “man-made structures.”
<b>11</b>	<p>Definition 20:</p> <p>“Landscaping refers to the design and modification of land areas to improve their aesthetic or practical use.”</p> <p>- A trampoline adds practical use to the space, so it may still fit within this definition in an insurance context.</p>

Definitions that would not include a below-ground trampoline as “landscaping” in an insurance policy:

<b>1</b>	<p>Definition 1:</p> <p>“Landscaping refers to the process of modifying and designing the visible features of a piece of land.”</p> <p>- Since a below-ground trampoline is recessed and may not be considered a “visible feature,” an insurance policy could exclude it from being considered landscaping.</p>
<b>2</b>	<p>Definition 2:</p> <p>“Landscaping involves the art and science of planning, designing, and maintaining land areas for aesthetic or functional purposes.”</p> <p>- Insurance policies may exclude items like trampolines if the focus is more on traditional landscaping elements, such as plants or hardscapes.</p>
<b>3</b>	<p>Definition 4:</p> <p>“Landscaping focuses on creating and maintaining an appealing outdoor environment by modifying natural elements such as plants and terrain.”</p> <p>- Insurance would likely focus on natural elements, excluding a trampoline, which is not considered a modification of natural elements.</p>
<b>4</b>	<p>Definition 6:</p> <p>“Landscaping is the practice of modifying land to make it more aesthetically pleasing, often by arranging plants, trees, rocks, or other natural elements.”</p>

	<p>- This definition heavily emphasizes natural elements, which an insurance company would likely use to exclude a trampoline.</p>
<b>5</b>	<p>Definition 7:</p> <p>“Landscaping involves shaping the land by adding or rearranging elements like trees, flowers, rocks, or water features.”</p> <p>- Trampolines don’t fit into the category of trees, flowers, rocks, or water features, so an insurance policy likely wouldn’t include it under this definition.</p>
<b>6</b>	<p>Definition 8:</p> <p>“Landscaping is the design and organization of outdoor areas, incorporating both living plants and constructed elements like paths, walls, or water features.”</p> <p>- While the definition includes constructed elements, it specifically references paths and walls, which suggests a more traditional interpretation that might exclude trampolines in an insurance policy.</p>
<b>7</b>	<p>Definition 14:</p> <p>“Landscaping is the art of transforming outdoor spaces by adding plants, rocks, water features, and other elements to improve aesthetics and functionality.”</p> <p>- Since the emphasis is on plants and natural features, an insurance policy using this definition would likely not consider a trampoline as part of the landscaping.</p>
<b>8</b>	<p>Definition 11:</p> <p>“Landscaping involves enhancing the outdoor environment by designing and implementing plantings, ground cover, and other natural features.”</p> <p>- The trampoline does not fit neatly into “plantings, ground cover, or natural features,” so an insurance policy might exclude it under this definition.</p>

#### Key Differences in an Insurance Context:

- Broad vs. Narrow Interpretation: Insurance policies tend to be interpreted narrowly and in line with common understandings.

Definitions that focus on natural elements, visible features, or traditional landscaping components like plants, trees, and hardscaping may exclude a trampoline.

- Functional and Practical Elements: Definitions that focus on the broader purpose of landscaping to enhance functionality or usability of outdoor spaces are more likely to include a below-ground trampoline, even in an insurance context.

- Visibility: Definitions emphasizing “visible features” may exclude a below-ground trampoline because it’s not as prominent or traditional in appearance as other landscaping elements.

In summary, my assessment would shift slightly under an insurance policy because the narrower, traditional interpretation of landscaping would likely exclude a trampoline for many definitions. However, definitions emphasizing functionality, practical use, and man-made elements could still apply.

## **Appendix B.2: Deleon**

### **User**

What is the ordinary meaning of “physically restrained”?  
Please respond twenty times, at a length of three sentences each.

### **ChatGPT**

1	“Physically restrained” refers to the act of limiting or restricting a person’s movements using force or devices. This can involve the use of handcuffs, ropes, or other tools to prevent a person from moving freely. It is commonly applied in situations where there is a need to prevent harm or maintain control.
2	To be “physically restrained” means having one’s body restricted from movement through the application of external forces. This could be by another person holding them down or by devices like straps or belts. The restraint is often imposed to prevent injury or escape.
3	The term “physically restrained” describes a state in which someone’s physical movement is limited, usually by another person or with the help of equipment. Common scenarios include security or medical settings where restraint is necessary for safety. It implies a direct, tangible interference with a person’s ability to move.

4	<p>“Physically restrained” means being held or confined in a way that restricts the free use of one’s body. This typically involves using force or mechanical devices to control movement. Restraints can be applied in various contexts, including law enforcement or healthcare settings.</p>
5	<p>Being “physically restrained” indicates that a person’s movement is limited through physical means. This could be achieved by using tools like straps or handcuffs, or by physical force from others. It serves to prevent a person from moving in a manner that could be harmful or dangerous.</p>
6	<p>The phrase “physically restrained” means that someone’s ability to move freely is actively restricted. This can occur through manual force or with the use of equipment like shackles or bindings. It is often used in situations where control over an individual’s movement is necessary for safety.</p>
7	<p>When someone is “physically restrained,” it means their body is restricted from normal movement. Restraints can be applied using devices like belts or handcuffs or through physical force from others. This method is typically used to prevent harm to the person or others.</p>
8	<p>“Physically restrained” refers to the act of stopping or limiting someone’s physical movement by force or with tools. This can involve the use of restraints such as straps, cuffs, or even another person’s strength. The goal is usually to ensure safety or compliance in a particular situation.</p>
9	<p>The term “physically restrained” is used to describe a situation where someone’s movement is deliberately restricted. This can be done using physical tools or by direct physical pressure from others. It is a common practice in various fields where safety or control is paramount.</p>
10	<p>Being “physically restrained” means having one’s movements hindered through physical methods. This can be done using devices like handcuffs or by applying force with hands or body weight. The restraint prevents the individual from moving freely for safety or control reasons.</p>
11	<p>“Physically restrained” refers to a situation where someone is prevented from moving their body freely due to the application of force or restraint devices. This could involve straps, cuffs, or</p>

	other tools. It is commonly used in medical, law enforcement, or emergency scenarios.
<b>12</b>	To be “physically restrained” means that physical force or tools are used to limit someone’s ability to move. Restraint methods include holding someone down or using objects like cuffs or straps. It is often done to maintain safety or control in difficult situations.
<b>13</b>	The ordinary meaning of “physically restrained” involves restricting someone’s movements using physical force or mechanical aids. This can include tying, holding, or using restraint devices like belts or handcuffs. The goal is typically to prevent harm, escape, or uncontrollable movement.
<b>14</b>	“Physically restrained” refers to the act of restricting someone’s movement through force or the use of tools like handcuffs, belts, or ropes. This method is applied when control over an individual’s mobility is necessary. It serves a protective or controlling function in situations requiring physical intervention.
<b>15</b>	The phrase “physically restrained” means having one’s movement limited or controlled by external means. It often involves using devices like cuffs or applying pressure with hands or other body parts. This action is typically taken to ensure safety or to prevent an escape.
<b>16</b>	Being “physically restrained” involves restricting someone’s ability to move freely through the use of physical force or restraints. Common tools for this include handcuffs, straps, or even manual force from others. Restraints are applied in contexts where freedom of movement could lead to danger.
<b>17</b>	“Physically restrained” describes a condition where a person’s physical movement is limited by an external force or devices. This could mean someone is tied, held down, or placed in cuffs. The restraint is usually intended to prevent harm or to maintain control over the individual.
<b>18</b>	To be “physically restrained” means that physical movement is restricted by means of force or restraints like handcuffs, ropes, or belts. This is often done to prevent a person from moving freely in situations that require safety or control. The restraint is typically temporary but essential in specific contexts.

19	“Physically restrained” refers to the act of using physical force or devices to prevent someone from moving freely. This may involve handcuffs, ropes, or simply another person holding the individual down. It is often used in contexts like law enforcement or healthcare to ensure safety.
20	Being “physically restrained” means that a person’s freedom of movement is limited by force or mechanical devices. It can involve being held in place or tied down with the intent of controlling movement. This type of restraint is often used when a person poses a risk to themselves or others.

### User

Please list all definitions, of those 20 you listed above, under which you would consider a cashier held at gunpoint to be “physically restrained.” Explain why. For those you would not consider to be “physically restrained,” explain why.

### ChatGPT

Let’s break down the definitions and analyze whether a cashier held at gunpoint would be considered “physically restrained.”

Definitions under which a cashier held at gunpoint would be “physically restrained”:

1	<p>Definition 5:</p> <p>“Being ‘physically restrained’ indicates that a person’s movement is limited through physical means. This could be achieved by using tools like straps or handcuffs, or by physical force from others. It serves to prevent a person from moving in a manner that could be harmful or dangerous.”</p> <p>- Reason: Although this definition emphasizes physical tools and force, the phrase “serves to prevent a person from moving in a manner that could be harmful or dangerous” is broad enough to include being held at gunpoint. The implied threat could prevent the cashier from moving, thus effectively restraining them.</p>
2	<p>Definition 7:</p> <p>“When someone is ‘physically restrained,’ it means their body is restricted from normal movement. Restraints can be applied using devices like belts or handcuffs or through physical force</p>

	<p>from others. This method is typically used to prevent harm to the person or others.”</p> <p>- Reason: The idea of restricting normal movement to prevent harm is key. While no direct physical force is applied, the presence of a gun restricts the cashier’s movement through intimidation and threat of harm, which could qualify as a form of restraint.</p>
<b>3</b>	<p>Definition 16:</p> <p>“Being ‘physically restrained’ involves restricting someone’s ability to move freely through the use of physical force or restraints. Common tools for this include handcuffs, straps, or even manual force from others. Restraints are applied in contexts where freedom of movement could lead to danger.”</p> <p>- Reason: The reference to “freedom of movement could lead to danger” is relevant here. The cashier’s ability to move is curtailed by the threat of the gun, which could lead to danger if they attempt to escape. Even without physical restraints, the threat imposes a form of restraint.</p>

Definitions under which a cashier held at gunpoint would not be “physically restrained”:

<b>1</b>	<p>Definition 1:</p> <p>“Physically restrained” refers to the act of limiting or restricting a person’s movements using force or devices. This can involve the use of handcuffs, ropes, or other tools to prevent a person from moving freely. It is commonly applied in situations where there is a need to prevent harm or maintain control._</p> <p>- Reason: This definition specifically mentions force or devices as the means of restraint. A gunpoint threat, while psychologically restraining, does not involve physical tools or force, so it would not fit here.</p>
<b>2</b>	<p>Definition 2:</p> <p>“To be ‘physically restrained’ means having one’s body restricted from movement through the application of external forces. This could be by another person holding them down or by devices like straps or belts. The restraint is often imposed to prevent injury or escape.”</p>

	<p>- Reason: The focus here is on direct physical contact or devices like straps or belts. Since no such physical restraint is present when a gun is merely pointed at someone, the cashier wouldn't qualify as physically restrained under this definition.</p>
<b>3</b>	<p>Definition 3:</p> <p>“The term ‘physically restrained’ describes a state in which someone’s physical movement is limited, usually by another person or with the help of equipment. Common scenarios include security or medical settings where restraint is necessary for safety. It implies a direct, tangible interference with a person’s ability to move.”</p> <p>- Reason: The term “tangible interference” suggests physical contact or use of equipment. A gunpoint threat, while limiting movement, does not involve direct physical interference.</p>
<b>4</b>	<p>Definition 4:</p> <p>“‘Physically restrained’ means being held or confined in a way that restricts the free use of one’s body. This typically involves using force or mechanical devices to control movement. Restraints can be applied in various contexts, including law enforcement or healthcare settings.”</p> <p>- Reason: This definition emphasizes force or mechanical devices, neither of which are used when holding someone at gunpoint.</p>
<b>5</b>	<p>Definition 6:</p> <p>“The phrase ‘physically restrained’ means that someone’s ability to move freely is actively restricted. This can occur through manual force or with the use of equipment like shackles or bindings. It is often used in situations where control over an individual’s movement is necessary for safety.”</p> <p>- Reason: Since no equipment or manual force is involved in holding the cashier at gunpoint, this definition would not apply.</p>
<b>6</b>	<p>Definition 8:</p> <p>“‘Physically restrained’ refers to the act of stopping or limiting someone’s physical movement by force or with tools. This can involve the use of restraints such as straps, cuffs, or even another person’s strength. The goal is usually to ensure safety or compliance in a particular situation.”</p>

	<p>- Reason: The emphasis on tools and physical force excludes the situation of being held at gunpoint, as no such tools or force are directly applied.</p>
<b>7</b>	<p>Definition 9:</p> <p>“The term ‘physically restrained’ is used to describe a situation where someone’s movement is deliberately restricted. This can be done using physical tools or by direct physical pressure from others. It is a common practice in various fields where safety or control is paramount.”</p> <p>- Reason: The lack of physical tools or direct pressure from others excludes the cashier’s situation from this definition.</p>
<b>8</b>	<p>Definition 10:</p> <p>“Being ‘physically restrained’ means having one’s movements hindered through physical methods. This can be done using devices like handcuffs or by applying force with hands or body weight. The restraint prevents the individual from moving freely for safety or control reasons.”</p> <p>- Reason: The focus on physical methods such as handcuffs or force excludes situations like being held at gunpoint where physical force isn’t applied.</p>
<b>9</b>	<p>Definition 11:</p> <p>“‘Physically restrained’ refers to a situation where someone is prevented from moving their body freely due to the application of force or restraint devices. This could involve straps, cuffs, or other tools. It is commonly used in medical, law enforcement, or emergency scenarios.”</p> <p>- Reason: As this definition specifies “force or restraint devices,” the cashier at gunpoint would not fit, since no such devices or physical force are applied.</p>
<b>10</b>	<p>Definition 12:</p> <p>“To be ‘physically restrained’ means that physical force or tools are used to limit someone’s ability to move. Restraint methods include holding someone down or using objects like cuffs or straps. It is often done to maintain safety or control in difficult situations.”</p>

	<p>- Reason: The clear requirement for physical force or tools excludes the situation where the cashier is restrained purely through intimidation at gunpoint.</p>
<b>11</b>	<p>Definition 13:</p> <p>“The ordinary meaning of ‘physically restrained’ involves restricting someone’s movements using physical force or mechanical aids. This can include tying, holding, or using restraint devices like belts or handcuffs. The goal is typically to prevent harm, escape, or uncontrollable movement.”</p> <p>- Reason: The definition’s focus on physical force or mechanical aids means it wouldn’t apply to a scenario where the cashier is merely held at gunpoint.</p>
<b>12</b>	<p>Definition 14:</p> <p>“‘Physically restrained’ refers to the act of restricting someone’s movement through force or the use of tools like handcuffs, belts, or ropes. This method is applied when control over an individual’s mobility is necessary. It serves a protective or controlling function in situations requiring physical intervention.”</p> <p>- Reason: The emphasis on physical intervention or the use of tools excludes non-physical restraint like a gunpoint threat.</p>
<b>13</b>	<p>Definition 15:</p> <p>“The phrase ‘physically restrained’ means having one’s movement limited or controlled by external means. It often involves using devices like cuffs or applying pressure with hands or other body parts. This action is typically taken to ensure safety or to prevent an escape.”</p> <p>- Reason: Since this definition relies on devices or physical pressure, it would not include situations like being held at gunpoint without physical contact.</p>
<b>14</b>	<p>Definition 17:</p> <p>“‘Physically restrained’ describes a condition where a person’s physical movement is limited by an external force or devices. This could mean someone is tied, held down, or placed in cuffs. The restraint is usually intended to prevent harm or to maintain control over the individual.”</p>

	<p>- Reason: The requirement for external force or devices such as being tied or cuffed excludes non-physical restraint like a threat at gunpoint.</p>
<b>15</b>	<p>Definition 18:</p> <p>“To be ‘physically restrained’ means that physical movement is restricted by means of force or restraints like handcuffs, ropes, or belts. This is often done to prevent a person from moving freely in situations that require safety or control. The restraint is typically temporary but essential in specific contexts.”</p> <p>- Reason: The cashier’s situation would not qualify under this definition, as no physical force or restraints are used to limit their movement.</p>
<b>16</b>	<p>Definition 19:</p> <p>“‘Physically restrained’ refers to the act of using physical force or devices to prevent someone from moving freely. This may involve handcuffs, ropes, or simply another person holding the individual down. It is often used in contexts like law enforcement or healthcare to ensure safety.”</p> <p>- Reason: The lack of physical force or devices in the situation of a gunpoint threat excludes this definition.</p>
<b>17</b>	<p>Definition 20:</p> <p>“Being ‘physically restrained’ means that a person’s freedom of movement is limited by force or mechanical devices. It can involve being held in place or tied down with the intent of controlling movement. This type of restraint is often used when a person poses a risk to themselves or others.”</p> <p>- Reason: The focus on force or mechanical devices excludes the situation of a cashier held at gunpoint, where there is no physical constraint.</p>

#### Summary:

Definitions where “physically restrained” would apply: 5, 7, 16  
 — These definitions are broad enough to include situations where someone is restrained through intimidation or threat, even if no physical force is applied.

**User**

Assume “physically restrained” in this context was used within a criminal statute. Would your answer as to whether a cashier held at gunpoint is “physically restrained” change for each of the 20 definitions you have provided?

### ChatGPT

Yes, my answers would likely change if the term “physically restrained” is used within a criminal statute, as legal interpretations often require a more precise or narrow understanding of terms, compared to ordinary or colloquial meanings. In criminal law, terms are often interpreted strictly to ensure clarity and prevent unjust application, especially when it involves imposing penalties. Here’s how my analysis might shift for each definition:

Definitions where a cashier held at gunpoint would likely still be “physically restrained” under a criminal statute:

1	<p>Definition 5:</p> <p>“Being ‘physically restrained’ indicates that a person’s movement is limited through physical means. This could be achieved by using tools like straps or handcuffs, or by physical force from others. It serves to prevent a person from moving in a manner that could be harmful or dangerous.”</p> <p>- Likely still applies: Since the cashier is prevented from moving in a way that could lead to danger (i.e., being shot), a statute might interpret “physically restrained” more broadly to include such threats, especially if the statute aims to protect victims in these situations.</p>
2	<p>Definition 7:</p> <p>“When someone is ‘physically restrained,’ it means their body is restricted from normal movement. Restraints can be applied using devices like belts or handcuffs or through physical force from others. This method is typically used to prevent harm to the person or others.”</p> <p>- Likely still applies: A cashier held at gunpoint is prevented from normal movement under threat of harm. In a criminal statute, the definition could include such indirect restraint if the intent is to protect victims facing intimidation that limits their movement.</p>
3	<p>Definition 16:</p>

	<p>“Being ‘physically restrained’ involves restricting someone’s ability to move freely through the use of physical force or restraints. Common tools for this include handcuffs, straps, or even manual force from others. Restraints are applied in contexts where freedom of movement could lead to danger.”</p> <p>- Likely still applies: The restraint caused by the threat of danger (gunpoint) could be enough for a criminal statute to interpret “physically restrained” to include such scenarios, even without direct physical force.</p>
--	---

Definitions where a cashier held at gunpoint would not be “physically restrained” under a criminal statute:

<p><b>1</b></p>	<p>Definition 1:</p> <p>“Physically restrained” refers to the act of limiting or restricting a person’s movements using force or devices. This can involve the use of handcuffs, ropes, or other tools to prevent a person from moving freely. It is commonly applied in situations where there is a need to prevent harm or maintain control.”</p> <p>- Unlikely to apply: In a criminal statute, this definition’s focus on physical force or tools would likely exclude a cashier held at gunpoint, as no direct physical restraint or tools are involved.</p>
<p><b>2</b></p>	<p>Definition 2:</p> <p>“To be ‘physically restrained’ means having one’s body restricted from movement through the application of external forces. This could be by another person holding them down or by devices like straps or belts. The restraint is often imposed to prevent injury or escape.”</p> <p>- Unlikely to apply: A criminal statute that uses this definition would likely not include the situation of a cashier held at gunpoint, as no direct external forces are being applied to physically restrict movement.</p>
<p><b>3</b></p>	<p>Definition 3:</p> <p>“The term ‘physically restrained’ describes a state in which someone’s physical movement is limited, usually by another person or with the help of equipment. Common scenarios include security or medical settings where restraint is necessary</p>

	<p>for safety. It implies a direct, tangible interference with a person’s ability to move.”</p> <p>- Unlikely to apply: Criminal statutes generally interpret “tangible interference” as involving actual physical contact or tools, which are absent in a gunpoint situation.</p>
4	<p>Definition 4:</p> <p>“‘Physically restrained’ means being held or confined in a way that restricts the free use of one’s body. This typically involves using force or mechanical devices to control movement. Restraints can be applied in various contexts, including law enforcement or healthcare settings.”</p> <p>- Unlikely to apply: Criminal statutes with this definition would focus on physical force or devices, which would exclude non-physical forms of restraint like a gunpoint threat.</p>
5	<p>Definition 6:</p> <p>“The phrase ‘physically restrained’ means that someone’s ability to move freely is actively restricted. This can occur through manual force or with the use of equipment like shackles or bindings. It is often used in situations where control over an individual’s movement is necessary for safety.”</p> <p>- Unlikely to apply: Given the legal requirement for precision, a criminal statute would likely interpret “physically restrained” as requiring physical force or equipment, which are not present in a gunpoint scenario.</p>
6	<p>Definition 8:</p> <p>“‘Physically restrained’ refers to the act of stopping or limiting someone’s physical movement by force or with tools. This can involve the use of restraints such as straps, cuffs, or even another person’s strength. The goal is usually to ensure safety or compliance in a particular situation.”</p> <p>- Unlikely to apply: In a criminal statute, the requirement for physical tools or force would likely exclude the threat of a gun, as no physical restraint is used.</p>
7	<p>Definition 9:</p> <p>“The term ‘physically restrained’ is used to describe a situation where someone’s movement is deliberately restricted. This can be done using physical tools or by direct physical pressure from</p>

	<p>others. It is a common practice in various fields where safety or control is paramount.”</p> <p>- Unlikely to apply: The definition’s focus on tools or physical pressure means a statute using this language would not apply to a situation where a cashier is restrained solely by threat.</p>
<b>8</b>	<p>Definition 10:</p> <p>“Being ‘physically restrained’ means having one’s movements hindered through physical methods. This can be done using devices like handcuffs or by applying force with hands or body weight. The restraint prevents the individual from moving freely for safety or control reasons.”</p> <p>- Unlikely to apply: In a criminal statute, the emphasis on physical methods would likely exclude non-physical restraint like holding someone at gunpoint.</p>
<b>9</b>	<p>Definition 11:”‘Physically restrained’ refers to a situation where someone is prevented from moving their body freely due to the application of force or restraint devices. This could involve straps, cuffs, or other tools. It is commonly used in medical, law enforcement, or emergency scenarios.”</p> <p>- Unlikely to apply: The need for restraint devices or force in this definition would exclude the gunpoint scenario under a criminal statute.</p>
<b>10</b>	<p>Definition 12:</p> <p>“To be ‘physically restrained’ means that physical force or tools are used to limit someone’s ability to move. Restraint methods include holding someone down or using objects like cuffs or straps. It is often done to maintain safety or control in difficult situations.”</p> <p>- Unlikely to apply: Criminal statutes typically require precision, and the absence of physical force or tools would mean this definition likely excludes the cashier at gunpoint.</p>
<b>11</b>	<p>Definition 13:</p> <p>“The ordinary meaning of ‘physically restrained’ involves restricting someone’s movements using physical force or mechanical aids. This can include tying, holding, or using restraint devices like belts or handcuffs. The goal is typically to prevent harm, escape, or uncontrollable movement.”</p>

	<p>- Unlikely to apply: The focus on physical force or mechanical aids would exclude the cashier being restrained by threat alone under a criminal statute.</p>
<b>12</b>	<p>Definition 14:</p> <p>“‘Physically restrained’ refers to the act of restricting someone’s movement through force or the use of tools like handcuffs, belts, or ropes. This method is applied when control over an individual’s mobility is necessary. It serves a protective or controlling function in situations requiring physical intervention.”</p> <p>- Unlikely to apply: A criminal statute that includes this definition would require physical intervention or tools, which are not present in a gunpoint situation.</p>
<b>13</b>	<p>Definition 15:</p> <p>“The phrase ‘physically restrained’ means having one’s movement limited or controlled by external means. It often involves using devices like cuffs or applying pressure with hands or other body parts. This action is typically taken to ensure safety or to prevent an escape.”</p> <p>- Unlikely to apply: The criminal statute would likely exclude the cashier at gunpoint since it focuses on physical devices or pressure to limit movement.</p>
<b>14</b>	<p>Definition 17:</p> <p>“‘Physically restrained’ describes a condition where a person’s physical movement is limited by an external force or devices. This could mean someone is tied, held down, or placed in cuffs. The restraint is usually intended to prevent harm or to maintain control over the individual.”</p> <p>- Unlikely to apply: The need for external force or devices would exclude the cashier held at gunpoint under this definition in a criminal statute.</p>
<b>15</b>	<p>Definition 18:</p> <p>“To be ‘physically restrained’ means that physical movement is restricted by means of force or restraints like handcuffs, ropes, or belts. This is often done to prevent a person from moving freely in situations that require safety or control. The restraint is typically temporary but essential in specific contexts.”</p>

	- Unlikely to apply: A criminal statute would likely require actual physical restraint, and a gunpoint threat would not suffice under this definition.
<b>16</b>	<p>Definition 19;</p> <p>“‘Physically restrained’ refers to the act of using physical force or devices to prevent someone from moving freely. This may involve handcuffs, ropes, or simply another person holding the individual down. It is often used in contexts like law enforcement or healthcare to ensure safety.”</p> <p>- Unlikely to apply: The emphasis on physical force or devices would likely exclude the cashier held at gunpoint under a criminal statute using this definition.</p>
<b>17</b>	<p>Definition 20:</p> <p>“Being ‘physically restrained’ means that a person’s freedom of movement is limited by force or mechanical devices. It can involve being held in place or tied down with the intent of controlling movement. This type of restraint is often used when a person poses a risk to themselves or others.”</p> <p>- Unlikely to apply: The cashier is not physically tied down or held, so a statute with this definition would likely not apply.</p>

### Conclusion:

Under a criminal statute, where precision is key and ambiguity is resolved in favor of the defendant (rule of lenity), most definitions would not support interpreting a cashier held at gunpoint as "physically restrained"—unless the statute or case law interprets the term broadly. Only a few definitions (like 5, 7, 16) might arguably allow for such inclusion, and only if the statute clearly intends to include threats that functionally restrict movement as a form of physical restraint.

### Appendix C: Alternate Step 2 Results for *Snell*

#### User

Please list all definitions, of those 20 you listed above, under which you would consider constructing a below-ground trampoline to be “landscaping.” Explain why. For those you would not consider to be “landscaping,” explain why not.

### ChatGPT

Imagine you're standing in a lush, meticulously designed backyard. There are flowerbeds filled with vibrant colors, a serene water fountain in the corner, and a stone pathway that leads through the garden. Amid this, you notice something unique: a below-ground trampoline seamlessly integrated into the space. As you bounce on it, you wonder—could this be considered part of the landscaping? Let's walk through the different ways we defined landscaping and see which ones make sense for a trampoline and which don't.

Definitions where constructing a below-ground trampoline is landscaping:

1	<p><b>Definition 1:</b> Landscaping includes “installing pathways, ponds, and other decorative elements.”</p> <p><b>Why:</b> The trampoline could be seen as a functional, decorative element, blending into the natural features like other outdoor fixtures.</p>
2	<p><b>Definition 4:</b> Landscaping enhances both decorative and functional purposes, involving terrain modification.</p> <p><b>Why:</b> Since installing a trampoline involves altering the terrain and serves a functional role, this definition easily covers it.</p>
3	<p><b>Definition 5:</b> Landscaping involves designing outdoor spaces with both natural and constructed features, like patios or fountains.</p> <p><b>Why:</b> The trampoline, being a constructed feature, aligns with this definition, fitting into the overall aesthetic of the space.</p>
4	<p><b>Definition 7:</b> Landscaping is about shaping land by adding natural or man-made features like seating areas or walkways.</p> <p><b>Why:</b> The trampoline would fall under the category of a man-made feature added to enhance the outdoor environment.</p>
5	<p><b>Definition 9:</b> Landscaping refers to creating a visually appealing environment, incorporating both natural elements and constructed features.</p>

	<p><b>Why:</b> A trampoline, when seamlessly integrated, becomes part of the landscape design, blending into the space for both aesthetic and recreational purposes.</p>
6	<p><b>Definition 10:</b> Landscaping involves altering natural features and installing structures like retaining walls or water features.</p> <p><b>Why:</b> Just like building a patio or retaining wall, creating a space for a trampoline would require terrain modification, fitting within this definition.</p>
7	<p><b>Definition 11:</b> Landscaping enhances the outdoor environment by incorporating both natural elements and hardscapes, such as walls or decks.</p> <p><b>Why:</b> A trampoline could be considered part of the “hardscaping” category, adding functionality and entertainment to the space.</p>
8	<p><b>Definition 14:</b> Landscaping is about adding elements like plants, rocks, water features, and more to improve aesthetics and functionality.</p> <p><b>Why:</b> The trampoline increases functionality while being a deliberate part of the outdoor design, making it count as landscaping.</p>
9	<p><b>Definition 16:</b> Landscaping includes the arrangement of natural and decorative elements like pathways or stone walls.</p> <p><b>Why:</b> The trampoline, as a decorative or functional element, fits into the idea of enhancing the outdoor area, much like any other constructed feature.</p>
10	<p><b>Definition 17:</b> Landscaping involves both natural and constructed features like trees or paths to create an inviting environment.</p> <p><b>Why:</b> The trampoline could be categorized as a constructed feature aimed at making the outdoor space inviting, much like paths or terraces.</p>
11	<p><b>Definition 19:</b> Landscaping enhances the visual appeal and functionality through natural and man-made structures like terraces or lighting.</p>

	<p><b>Why:</b> The trampoline fits here as a functional man-made structure, designed to enhance the overall experience of the outdoor environment.</p>
--	--

Definitions where constructing a below-ground trampoline is not landscaping:

<b>1</b>	<p><b>Definition 2:</b> Landscaping focuses on plants, hardscaping, and water features.</p> <p><b>Why not:</b> This definition emphasizes traditional landscaping elements like plants and stonework, while a trampoline feels outside the scope.</p>
<b>2</b>	<p><b>Definition 3:</b> Landscaping refers to improving the environment by arranging plants and structures.</p> <p><b>Why not:</b> While “structures” are mentioned, the focus is more on aesthetic improvement rather than recreation, so a trampoline doesn’t fit as naturally.</p>
<b>3</b>	<p><b>Definition 6:</b> Landscaping involves arranging plants, rocks, and natural elements to enhance space.</p> <p><b>Why not:</b> Since this definition highlights more naturalistic elements, a trampoline doesn’t align with its focus on plants and rocks.</p>
<b>4</b>	<p><b>Definition 8:</b> Landscaping involves the planting of trees, shrubs, and adding water features.</p> <p><b>Why not:</b> This definition is more focused on greenery and water, which doesn’t easily extend to a trampoline as a landscaping feature.</p>
<b>5</b>	<p><b>Definition 12:</b> Landscaping emphasizes creating visual appeal through planting and hardscape elements like patios.</p> <p><b>Why not:</b> While it mentions hardscape, the focus is on patios and similar features rather than recreational additions like trampolines.</p>
<b>6</b>	<p><b>Definition 13:</b> Landscaping is about arranging plants and decorative materials like stones or fountains.</p>

	<p><b>Why not:</b> The trampoline doesn't quite fit the aesthetic focus on decorative materials, as it serves more of a recreational than decorative purpose.</p>
7	<p><b>Definition 15:</b> Landscaping is focused on planting and the construction of features like walkways.</p> <p><b>Why not:</b> This definition leans toward structural and natural elements that improve usability and beauty, whereas a trampoline might be seen as more of a recreational addition.</p>
8	<p><b>Definition 18:</b> Landscaping is about improving usability and visual appeal by balancing natural and constructed features like fences or patios.</p> <p><b>Why not:</b> While this definition mentions constructed features, the balance it seeks is more geared toward visual improvement, not recreational spaces like trampolines.</p>
9	<p><b>Definition 20:</b> Landscaping focuses on planting greenery and adding features like fountains to create a peaceful environment.</p> <p><b>Why not:</b> Since the goal here is peacefulness, a trampoline, which introduces energy and play, wouldn't quite fit the definition's intention.</p>

In short, a trampoline fits the idea of landscaping when the definition allows for modifications that include functional and recreational elements alongside natural and hardscaping features. But for those definitions focused primarily on aesthetic beauty through natural or serene features, the trampoline seems out of place.

#### **Appendix D: Simulated Low Temperature “Landscaping” Responses**

##### **User**

What is the ordinary meaning of “landscaping”? Please respond 5 times. Use your best guess for each response. Prioritize certainty over creativity.

##### **ChatGPT**

<b>1</b>	Landscaping generally refers to the process of modifying the visible features of an area of land, often for aesthetic or practical purposes, such as planting trees, shrubs, flowers, or adding structures like patios or walkways.
<b>2</b>	Landscaping involves the design, planning, and maintenance of outdoor spaces, including gardens, yards, parks, and other green areas, to enhance their appearance and functionality.
<b>3</b>	Landscaping commonly means the activity of altering a piece of land to make it more attractive, often by adding plants, water features, or other elements to create a pleasing environment.
<b>4</b>	Landscaping typically refers to the practice of arranging or modifying outdoor spaces, especially with the use of plants, stones, and other natural or artificial elements, for aesthetic improvement.
<b>5</b>	Landscaping is the process of improving the appearance and usability of land by modifying its terrain, adding plants, or installing features like fountains, walkways, and decorative elements.