

# YOUR PROMPT RESPONSE IS APPRECIATED: GENERATIVE AI AND THE FUTURE OF LAWYERS

*Matthew J. Tobin<sup>†</sup>*  
*Justin Lischak Earley<sup>‡</sup>*

## Abstract

Generative artificial intelligence (“GenAI”) took the world by storm in 2023. In its wake, some have speculated that GenAI is a serious threat to the livelihoods of lawyers, arguing that its capabilities will allow it to take on much of today’s legal work. In this paper, we suggest this superficial view misses the substantial economic opportunities GenAI may offer the legal profession. We will explain that fully utilizing GenAI requires deep substantive knowledge of a topic, and the ability to wield natural language with precision. The background we’ll provide on the inner workings of GenAI will show that its usefulness depends on how skillfully it’s “prompted” (asked to perform a task) and how thoughtfully its output is evaluated. We believe that the skills required to use GenAI effectively are the core skills of “thinking like a lawyer,” taught in law school and honed throughout years of legal practice. Therefore, rather than being a threat to lawyers, we propose that lawyers’ skills are an unrecognized key to unlocking the highest uses of GenAI. Indeed, GenAI may generate more lawyers, or at least more people who must learn to “think like a lawyer.”

---

<sup>†</sup> Matthew J. Tobin is Senior Corporate Technology Underwriter for First American Title Insurance Company. He received his law degree from the University of Wisconsin. Matthew can be found online at <https://www.linkedin.com/in/matthewjamestobin>

<sup>‡</sup> Justin Lischak Earley is Chief Innovation Underwriter for First American Title Insurance Company. Justin is a Fellow of the American College of Mortgage Attorneys and a Fellow of the American College of Real Estate Lawyers. He received his law degree from Duke University, and received his master’s degree in Human-Computer Interaction & Design from the University of California, Irvine. Justin can be found online at <https://jdlesq.com>

The views that we express in this article are our own personal scholarly opinions, and not necessarily the views of our employer. This is a purely academic article, and our views are not investment advice about any company (including but not limited to our employer), nor legal advice about anyone or any situation.

*“I’ve been talking to the main computer. It hates me.”*

*– Marvin, the Paranoid Android*<sup>1</sup>

**A**s lawyers, we are creatures of history. When we analyze questions, we do so by applying precedents to fact patterns. Many of these precedents are longstanding—perhaps centuries old. The facts change, yes, and all the time. But the law changes rarely. When change does occur, the process is methodical. In our daily lives we stand on the shoulders of many generations, doctrines, and black-letter rules that have come before us. Indeed, the very word “precedent” tells us what we look for to get answers: “what came before.”

But the pace at which generative artificial intelligence (“GenAI”) has advanced has worn out the word “unprecedented.” Since OpenAI’s product [ChatGPT](#) hit the public consciousness in March 2023, the arc of history has bent into a skyward parabola. While trying to keep pace with the speed of innovation surrounding an emerging technology is nothing new, the power of the internet and the very low barrier of entry to use GenAI has led to an explosion of resources. The rub, as usual, lies in knowing what information to trust. After all, people believed the iPhone was a doomed and useless product<sup>2</sup> and that the internet was nothing more than “baloney.”<sup>3</sup> The predictions around GenAI have been as varied as the human mind. To some, GenAI is too inconsistent to be useful; a mere

---

<sup>1</sup> THE HITCHHIKER’S GUIDE TO THE GALAXY (Touchstone Pictures et al. 2005).

<sup>2</sup> Seth Porges, *The Futurist: We Predict the iPhone Will Bomb*, TECHCRUNCH, June 7, 2007, <https://techcrunch.com/2007/06/07/the-futurist-we-predict-the-iphone-will-bomb/> (last visited Jan. 3, 2024) (“That virtual keyboard will be about as useful for tapping out emails and text messages as a rotary phone. Don’t be surprised if a sizable contingent of iPhone buyers express some remorse at ditching their BlackBerry when they spend an extra hour each day pumping out emails on the road.”).

<sup>3</sup> Clifford Stoll, *Why the Web Won’t be Nirvana*, NEWSWEEK, February 26, 1995, <https://www.newsweek.com/clifford-stoll-why-web-wont-be-nirvana-185306> (last visited Jan. 5, 2024) (“We’re promised instant catalog shopping—just point and click for great deals. We’ll order airline tickets over the network, make restaurant reservations and negotiate sales contracts. Stores will become obsolete. So how come my local mall does more business in an afternoon than the entire Internet handles in a month? Even if there were a trustworthy way to send money over the Internet—which there isn’t—the network is missing a most essential ingredient of capitalism: salespeople.”).

toy.<sup>4</sup> To others, it represents a seismic change for humanity which will lead to the mass replacement of jobs,<sup>5</sup> lawyers included.<sup>6</sup>

Looking for answers by reference to the companies at the forefront of this technology yields no clear insight. They seem equally prone to massive and rapid changes. During the writing of this paper alone we saw a boardroom coup of Napoleonic proportions at OpenAI (where the exile to Elba lasted only a weekend),<sup>7</sup> the launch of multiple new GenAI models by other tech companies ([Gemini](#) from Google, [Claude](#) from Anthropic, etc.), and speculation that OpenAI might have actually created “true” artificial general intelligence (“AGI”), under an approach mysteriously referred to as Q\* [q star].<sup>8</sup> At this pace, what can two real estate lawyers—creatures of history as we are—write about this subject that will still be valid by the time you are reading it?

A lot, we hope. While the pace of this technology’s advancement *is* unprecedented, we need not simply throw up our hands. At an earlier time of rapid innovation in technology, a strategist faced with a similarly uncertain future wrote:

[There may not be] *precedents* to be followed. But a precedent is different from and less valuable than a principle. The former may be originally faulty, or may cease to apply through change of circumstances; the latter has its

---

<sup>4</sup> Ian Bogost, *ChatGPT Is Dumber Than You Think*, THE ATLANTIC, December 7, 2022, <https://www.theatlantic.com/technology/archive/2022/12/chatgpt-openai-artificial-intelligence-writing-ethics/672386/> (last visited Jan. 5, 2024) (“Treat it like a toy, not a tool.”).

<sup>5</sup> *E.g.*, *Pause Giant AI Experiments: An Open Letter*, FUTURE OF LIFE INSTITUTE, Mar. 22, 2023, <https://futureoflife.org/open-letter/pause-giant-ai-experiments/> (last visited Jan. 3, 2024) (“Contemporary AI systems are now becoming human-competitive at general tasks, and we must ask ourselves . . . *Should* we automate away all the jobs, including the fulfilling ones? *Should* we develop nonhuman minds that might eventually outnumber, outsmart, obsolete and replace us?”) (emphasis in original).

<sup>6</sup> *E.g.*, <https://willrobotstakemyjob.com/lawyers> (last visited Jan. 3, 2024) (calculating a 22% chance that lawyers will be automated out of existence); *see also* Zack Womack, *1 in 4 Large Law Firms Expect Generative AI to Replace Jobs Internally in Next 5 Years*, LAW.COM, Aug. 21, 2023, <https://www.law.com/international-edition/2023/08/21/1-in-4-large-law-firms-expect-generative-ai-to-replace-jobs-internally-in-next-5-years/> (last visited Jan. 3, 2024) (“At least one in four top commercial law firms expect generative artificial intelligence to replace jobs within their firm within the next five years.”).

<sup>7</sup> *See, e.g.*, Lauren Goode & Will Knight, *Sam Altman to Return as CEO of OpenAI*, WIRED, Nov. 22, 2023, <https://www.wired.com/story/sam-altman-openai-back/> (last visited Jan. 3, 2024).

<sup>8</sup> *See, e.g.*, Will Knight, *These Clues Hint at the True Nature of OpenAI’s Shadowy Q\* Project*, WIRED, Nov. 30, 2023, <https://www.wired.com/story/fast-forward-clues-hint-openai-shadowy-q-project/> (last visited Jan. 3, 2024) (“the pace of development alarmed some researchers”).

root in the essential nature of things, and, however various its application as conditions change, remains a standard to which action must conform to attain success.<sup>9</sup>

In this paper, we take a principle-based approach. We look carefully at the essential nature of GenAI to consider how it may affect the lives of lawyers (specifically, real estate lawyers). We see basic principles behind GenAI that will stand the test of time. We lay down those principles as we see them from our vantage point, writing in January 2024. Of course, we may well be wrong. It has been reported that the foundational corporate documents of the OpenAI-Microsoft partnership contain an “all bets are off” clause if OpenAI should somehow create “true” artificial general intelligence that exceeds human capabilities.<sup>10</sup> We claim the same “out” here.

But barring the creation of [Skynet](#) or [Lt. Commander Data](#), we suggest that any hyperbole around GenAI heralding the end of lawyers is misplaced. Indeed, we propose the opposite: the skills needed to leverage GenAI to its fullest potential are those that *belong to lawyers*. This is because, as the leading voices in the field have described, using GenAI is “programming in natural language.”<sup>11</sup> And as one of us foreshadowed several years ago,<sup>12</sup> the skills it takes to code in natural language are those *of lawyers*: the ability to wield words with precision, a deep subject-matter command of those words’ meanings, and a firm grasp of the consequences that can flow from those words.

As we explain below, GenAI lacks those skills—they must be supplied by you, the user. To wit, if you have dabbled with ChatGPT, Bard, or any of the other commercially available GenAI models, you may find yourself agreeing with Marvin the Paranoid

---

<sup>9</sup> A.T. MAHAN, THE INFLUENCE OF SEA POWER UPON HISTORY, 1660–1783, *reprinted in* ROOTS OF STRATEGY, BOOK 4, at 62 (Stackpole Books 1999, ed. David Jablonsky).

<sup>10</sup> Steven Levy, *The Transformers*, WIRED, Oct. 2023, at 45 (Microsoft CEO Satya Nadella: “[AGI] might be the last invention of humanity . . . so we might have bigger issues to consider once machines are smarter than we are.”).

<sup>11</sup> See, e.g., Laria Reynolds & Kyle McDonell, *Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm*, 2021 EXTENDED ABSTRACTS OF THE CHI CONFERENCE ON HUMAN FACTORS IN COMPUTING SYSTEMS (May 2021), available at <https://dl.acm.org/doi/10.1145/3411763.3451760> (last visited Jan. 13, 2024).

<sup>12</sup> Justin Lischak Earley, *Refactoring Our Documents*, NEWS & NOTES, AMERICAN COLLEGE OF REAL ESTATE LAWYERS, Fall 2020, at 1, [https://jdl.esq.com/documents/Refactoring%20Our%20Documents%20\(14%20July%202020\).pdf](https://jdl.esq.com/documents/Refactoring%20Our%20Documents%20(14%20July%202020).pdf) (last visited Jan. 3, 2024) (“Indeed, there are good reasons to draw a parallel between ‘law coders’ and ‘software coders.’”).

Android’s sentiments quoted above: the computer hates you. The responses from GenAI models are bright, cheery, and often wrong or outright fabrications. However, as software developers often say, “that’s not a bug, that’s a feature!” Those commercial GenAI models are trained to be bright, cheery, and as *helpful as possible*. In general, they will answer a question asked by confidently replying with *something*.

But GenAI models have trouble knowing what they know. Perhaps because some attorneys were unaware of this, they forgot the old adage, “trust but verify.” Several high-profile and embarrassing filings submitted with fictional cases provided by a “helpful” GenAI model now serve as warnings to the legal community.<sup>13</sup> While these [double-facepalm](#) moments effectively highlight the risks from misusing GenAI or failing to maintain proper oversight, they should not be over-read. Any technology can be misused by those who fail to understand it. A handful of public failures should not overshadow the potential value that a nuanced, well-informed use of GenAI can deliver.

We need a clear-eyed view into what GenAI is, what it is not, and what the evidence shows best enables one to leverage its powers. We begin that journey by pointing out two fallacies about GenAI that cannot be allowed to cloud the reader’s judgment.

## 1. THE DANCING BEAR, THE PICKUP TRUCK, AND THE ANTHROPOMORPHIC FALLACY

*“Can machines think?”*

– Alan Turing <sup>14</sup>

Since time immemorial, people have been fascinated by creatures trained to perform human tasks. Among the most enduring motifs is that of “the dancing bear.”<sup>15</sup> Human beings tend to anthropomorphize the things around us. We want to believe that the things with which we interact are just like we are. We regularly anthropomorphize animals, and

---

<sup>13</sup> The latest suspected incident involves Michael Cohen’s attorney. *See, e.g.,* Ewan Palmer, *Michael Cohen Lawyer’s Citing of Fictional Cases Raises ChatGPT Suspicions*, NEWSWEEK, Dec. 14, 2023, <https://www.newsweek.com/michael-cohen-lawyer-chatgpt-court-cases-ai-1852530> (last visited Jan. 3, 2024).

<sup>14</sup> Alan M. Turing, *Computing Machinery and Intelligence*, 69 MIND 433, 433 (1950), available at <https://doi.org/10.1093/mind/LIX.236.433> (last visited Jan. 13, 2024).

<sup>15</sup> *Dancing Bears*, BEAR CONSERVATION, May 29, 2023, <http://www.bearconservation.org.uk/dancing-bears/> (last visited Jan. 3, 2024).

from our earliest conscious ages, we imagine that our toys have feelings, emotions, dreams. We envision their stories as part of our own. This is how a batch of straw becomes a doll, which in turn becomes a treasured childhood companion.<sup>16</sup>

When it comes to our technologies, we do much the same. We create things that look like us, and then imbue them with humanlike qualities that carry social meaning to us.<sup>17</sup> Consider the road the next time you drive. Pickup trucks “look” like burly, macho lumberjacks. Sports cars “look” like aggressive, competitive athletes. The designs we choose for these things mimic what we want ourselves to be.<sup>18</sup>

But the pickup truck is not a macho lumberjack. The sports car is not a competitive athlete. These tools are not self-aware. They have no feelings, emotions, or dreams. They are simply tools. Their human qualities are only inside our own heads.<sup>19</sup> To consider them as self-aware, thinking creatures is to fall into the anthropomorphic fallacy:

The anthropomorphic fallacy refers to the misleading attribution of human characteristics, emotions, motivations, or intentions to non-human entities, such as animals, objects, or natural phenomena.<sup>20</sup>

The anthropomorphic fallacy is what gives the dancing bear its enduring relevance. Dance is a human expressive art that can span the entire gamut of human emotional range. It drips with social meaning. A dance can be uplifting or depressing, aggressive or lethargic, salacious or monotonous. We are pre-wired to *want* to believe that the dancing bear must be expressing some work of art or emotion. It must be “like us.”

---

<sup>16</sup> *E.g.*, DON NORMAN, EMOTIONAL DESIGN: WHY WE LOVE (OR HATE) EVERYDAY THINGS 136 (Basic Books paperback ed. 2005) (“[A]nthropomorphism [is] the attribution of human motivations, beliefs, and feelings to animals and inanimate things. The more behavior something exhibits, the more we are apt to do this.”).

<sup>17</sup> *See, e.g., id.* at 138 (“Designers take note. Humans are predisposed to anthropomorphize, to project human emotions and beliefs into anything.”).

<sup>18</sup> *See, e.g., id.* at 87 (“Pretty products—sexy automobiles, powerful-looking trucks, seductive bottles for drinks and perfume—play with the visceral level.”).

<sup>19</sup> *E.g., id.* at 194 (“When machines display emotions, they provide a rich and satisfying interaction with people, even though most of . . . [all that] comes from within the head of the person, not from the artificial system.”).

<sup>20</sup> <https://psychology.tips/anthropomorphic-fallacy/> (last visited Jan. 3, 2024).

But the dancing bear is not like us. Bears in the wild do not salsa or waltz as part of their natural order. The captured bear “dances” because it was (cruelly) trained to do so:

Typically bears would be taught to dance by placing them onto platforms of metal above large piles of burning logs. As the metal became hot the bears would be forced onto their hind legs by the use of the pole and nose ring and would then begin lifting each paw in turn to relieve them from the heat. As the process continued a drum or other music was played which, over a number of weeks and months, the bear came to associate with the pain in its feet. Subsequently whenever the bear heard the drum or music it would begin to ‘dance.’<sup>21</sup>

In pointing out the anthropomorphic fallacy, we need not delve into questions of what it means to be “self-aware.”<sup>22</sup> We are neither neuroscientists nor philosophers, but we feel comfortable that the bear does not dance to express emotion or artistry. It dances because it was trained to. So also with GenAI.<sup>23</sup>

The anthropomorphic fallacy looms large over any discussions of GenAI. After all, GenAI models are generally rated by direct human evaluation or by comparing their outputs against human responses.<sup>24</sup> The companies behind any type of AI are thus motivated to have their models produce something indistinguishable from that of a

---

<sup>21</sup> *Dancing Bears*, *supra* n.15.

<sup>22</sup> These are deep and difficult questions that are far beyond our training and experience. *See, e.g.*, Grace Huckins, *Machines Like Us*, MIT TECHNOLOGY REVIEW, Nov. / Dec. 2023, at 33:

As slippery a topic as consciousness can be, it is not impossible to pin down—put as simply as possible, it’s the ability to experience things. It’s often confused with terms like ‘sentience’ and ‘self-awareness,’ but according to the definitions that many experts use, consciousness is a prerequisite for those other, more sophisticated abilities. To be sentient, a being must be able to have positive and negative experiences—in other words, pleasures and pains. And being self-aware means not only having an experience but *knowing* that you are having an experience.

<sup>23</sup> *See, e.g., id.* at 35 (“That’s where LLMs like LaMDA currently are: they don’t possess the right type of feedback connections, use global workspaces, or appear to have any other markers of consciousness.”).

<sup>24</sup> Kyle Orland, *Turing test on steroids: Chatbot Arena crowdsourcing ratings for 45 AI models*, ARS TECHNICA, Dec. 13, 2023, <https://arstechnica.com/ai/2023/12/turing-test-on-steroids-chatbot-arena-crowdsources-ratings-for-45-ai-models/> (last visited Jan. 6, 2024) (“For those looking for a more rigorous way of comparing various models, the folks over at the Large Model Systems Organization (LMSys) have set up Chatbot Arena, a platform for generating Elo-style rankings for LLMs based on a crowdsourced blind-testing website.”).

human being. Indeed, the original and enduring evaluation for AI is the Turing test, introduced by Alan Turing in 1950 as a thought experiment to answer the question of whether machines are capable of thinking.<sup>25</sup>

The Turing test is conducted by a human evaluator having two distinct text-based conversations. The evaluator is made aware that one of their conversation partners is human, while the other is a machine. A machine is said to have passed the Turing test if the evaluator cannot distinguish which participant is human.<sup>26</sup> Therefore, creators of machines attempting to pass the Turing test are not necessarily incentivized to provide *correct* answers, but instead to create the *most human-like* response.

Whether that test has been passed since its inception is a matter of some debate. The first machine that may have done so was “ELIZA,” in 1966. ELIZA acted like a therapist, replying only to certain words. If the conversation partner used specific words it knew, ELIZA transformed those words into full sentences and replied, fooling some participants.<sup>27</sup> If ELIZA was the therapist, PARRY, created in 1972 by a Stanford scientist, was the patient. PARRY was designed to talk like a paranoid schizophrenic. And yes, PARRY and ELIZA “talked” with each other, and the results were predictably insane.<sup>28</sup> Progress thereafter was slow. The next time machines arguably passed the Turing test was in 2014, when the chatbot “Eugene Goostman” fooled 1 out of 3 judges.<sup>29</sup>

GenAI changed this slow progress dramatically. In no small part, this is because the responses from GenAI models are far more human-like than any machine or tool that previously took the Turing test. Researchers in 2023 presented a modified version of the Turing test to 1.5 million users for GenAI, and those users made 10 million guesses as to whether their conversation partner was a chatbot or a human. The participants correctly identified their partner as GenAI only *slightly* better than by guessing.<sup>30</sup>

---

<sup>25</sup> Turing, *supra* n.14.

<sup>26</sup> *Cf. id.* at 433–34.

<sup>27</sup> *See, e.g.,* NORMAN, EMOTIONAL DESIGN, *supra* n.16, at 188-91 (describing ELIZA’s accidentally fooling a programmer’s boss into thinking the boss was talking to the programmer).

<sup>28</sup> Megan Garber, *When PARRY met ELIZA: A Ridiculous Chatbot Conversation from 1972*, THE ATLANTIC, June 9, 2014, <https://www.theatlantic.com/technology/archive/2014/06/when-parry-met-eliza-a-ridiculous-chatbot-conversation-from-1972/372428/> (last visited Jan. 6, 2024).

<sup>29</sup> *Id.*

<sup>30</sup> <https://www.ai21.com/blog/human-or-not-results> (last visited Jan. 6, 2024).



Today, AI companies seek not just to imitate human beings, but to surpass them. The race to create and ethically control true “Artificial General Intelligence” (“AGI”) is fierce. AGI is either defined as AI that can understand, learn, and apply its intelligence to a wide range of problems in a manner similar to human intelligence, or AI systems that are “generally smarter than humans.”<sup>31</sup> After discoveries in 2017,<sup>32</sup> the creation of GenAI models and machines significantly advanced towards these objectives.<sup>33</sup> With the introduction of ChatGPT 3.5 to the public in early 2023, the discussions about the pace of GenAI and whether it was (or was about to become) AGI began running wild.<sup>34</sup>

It is no wonder that some people are questioning whether GenAI has crossed some ill-defined line into self-awareness. A former Google engineer quit his job over his belief that a GenAI called “LaMDA” is self-aware.<sup>35</sup> A plaintiff in a landmark lawsuit argued that the GenAI he invented can be an “author” of an artwork.<sup>36</sup> Again, thankfully, we need not answer these weighty questions here.

Rather, we need only constantly remind ourselves that, like [Fox Mulder](#), *we want to believe*,<sup>37</sup> and we must recognize our instinctual bias. As the aphorism goes, “[t]he wonder isn’t that the bear dances well but that the bear dances at all.”<sup>38</sup> In that vein, the

---

<sup>31</sup> See <https://openai.com/about> (last visited Jan. 15, 2024).

<sup>32</sup> Ashish Vaswani *et al.*, *Attention is All You Need*, 31ST CONFERENCE ON NEURAL INFORMATION PROCESSING SYSTEMS (2017), available at <https://papers.neurips.cc/paper/7181-attention-is-all-you-need.pdf> (last visited Jan. 13, 2024).

<sup>33</sup> *Cf.*, e.g., Eduardo Muñoz, *Attention is All You Need: Discovering the Transformer Paper*, TOWARDS DATA SCIENCE, Nov. 2, 2020, <https://towardsdatascience.com/attention-is-all-you-need-discovering-the-transformer-paper-73e5ff5e0634> (last visited Jan. 13, 2024).

<sup>34</sup> E.g., Will Knight, *Some Glimpse AGI in ChatGPT. Others Call it a Mirage*, WIRED, Apr. 18, 2023, <https://www.wired.com/story/chatgpt-agi-intelligence/> (last visited Jan. 13, 2024).

<sup>35</sup> Steven Levy, *Blake Lemoine Says Google’s LaMDA AI Faces ‘Bigotry’*, WIRED, June 17, 2022, <https://www.wired.com/story/blake-lemoine-google-lambda-ai-bigotry/> (last visited Jan. 3, 2024).

<sup>36</sup> Franklin Graves, *Thaler Pursues Copyright Challenge Over Denial of AI-Generated Work Registration*, June 6, 2022, IP WATCHDOG, <https://ipwatchdog.com/2022/06/06/thaler-pursues-copyright-challenge-denial-ai-generated-work-registration/id=149463/> (last visited Jan. 3, 2024); COOLEY LLP, *District Court Confirms ‘Human Authorship’ Requirement, Sets Copyright Boundary for AI-Generated Works*, Aug. 24, 2023, <https://www.cooley.com/news/insight/2023/2023-08-24-district-court-confirms-human-authorship-requirement-sets-copyright-boundary-for-ai-generated-works/> (last visited Jan. 3, 2024).

<sup>37</sup> Ella Morton, *The X-Files ‘I Want to Believe’ Poster’s Origin Story*, THE NEW REPUBLIC, Dec. 29, 2015, <https://newrepublic.com/article/126715/x-files-i-want-believe-posters-origin-story> (last visited Jan. 3, 2024).

<sup>38</sup> E.g., ALAN COOPER, *THE INMATES ARE RUNNING THE ASYLUM: WHY HIGH-TECH PRODUCTS DRIVE US CRAZY AND HOW TO RESTORE THE SANITY* 26 (Sams Publishing 2004) (emphasis removed).

problem with GenAI is twofold: the first is that the bear dances *very well*; so well that it can be exceptionally hard for us to resist the temptation to believe. The second is the very human language used to discuss GenAI: “learning,” “training,” “chatting.” That language baits us into believing and acting as if these tools were actually human.<sup>39</sup>

Staying out of the anthropomorphic fallacy trap allows us to avoid the breathless hyperbole surrounding GenAI and to avoid misunderstanding its capabilities. It reminds us that GenAI is just a tool that people have created. As a leading philosopher put it:

[AI instances] are intelligent tools, not colleagues. Don’t think of them as colleagues, don’t try to make them colleagues, and above all, don’t kid yourself that they are colleagues.<sup>40</sup>

AIs are tools, just like pickup trucks. Humanity knows how pickup trucks are built. And although we do not know how our own consciousness works or whether a GenAI is conscious, we *do* know how GenAIs work: the transformer.

## 2. TRANSFORMERS, MATHEMATICS IN DISGUISE, AND THE ENCHANTED DETERMINISM FALLACY

*“Any sufficiently advanced technology is indistinguishable from magic.”*

– Arthur C. Clarke <sup>41</sup>

Although it may appear to be magic, GenAI is really just math.<sup>42</sup> (Which may seem equally mysterious to us lawyers.) Jokes aside, it is important not to be intimidated by the concepts behind GenAI. This is so for two reasons.

---

<sup>39</sup> See, e.g., NORMAN, EMOTIONAL DESIGN, *supra* n.16, at 189 (“You can see how you might get captured by the conversation: your concerns received sympathetic responses. But Eliza has no understanding of language. It simply finds patterns and responds appropriately . . .”).

<sup>40</sup> John Thornhill, *Philosopher Daniel Dennett on AI, Robots and Religion*, FINANCIAL TIMES, Mar. 2, 2017 (quotation from Dennett).

<sup>41</sup> *Clarke’s Three Laws*, WIKIPEDIA, [https://en.wikipedia.org/wiki/Clarke%27s\\_three\\_laws](https://en.wikipedia.org/wiki/Clarke%27s_three_laws) (last visited 3 Jan. 2024).

<sup>42</sup> E.g., *18 Tech Experts Discuss AI Myths That Should Be Debunked*, FORBES, Aug. 21, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/08/21/18-tech-experts-discuss-ai-myths-that-should-be-debunked/?sh=4e3d0c1b37c9> (last visited Jan. 3, 2024) (“AI isn’t magic; it’s just math.”).

First, those who wish to profit from technologies use the human yearning for the supernatural to their advantage, and people who don't understand how something works may be duped. We have seen firsthand how otherwise sharp and skeptical minds become wide-eyed and manipulable when presented with seemingly mystical AI capabilities. Two key voices critical of AI call this tendency "enchanted determinism":

AI systems are seen as enchanted, beyond the known world, yet deterministic in that they discover patterns that can be applied with predictive certainty to everyday life.<sup>43</sup>

Entrepreneurial technologists are aware of this human yearning, and they leverage it. A common approach here is the refrain, "no one really knows how the AI did this!", as though the AI is powered by mysterious, fundamental forces of the universe that are beyond our limited understanding.<sup>44</sup> In this way, they tap into the mystique expressed by Arthur C. Clarke's famous dictum quoted above.

The paradigm example is the infamous "Mechanical Turk." Built to impress European royalty, the Turk seemed to be an automaton that played chess against (and defeated) many great players of its time. Presented as a "thinking machine," it was in fact a hoax: hiding inside of its bowels was a human chess master who directed the machine's moves. Given the discussion above about the yearning for magic and enchanted determinism, it is no surprise that the machine was dressed in the garb of an Ottoman sorcerer.<sup>45</sup>

While it is true that human beings cannot yet explain the *precise* inner workings behind *specific* outputs from GenAIs, asking for this level of detail is unnecessary to make practical use of the technology. It could equally be said that we don't know which molecules of fuel interacted with which molecules of air in a pickup truck's internal

---

<sup>43</sup> KATE CRAWFORD, *ATLAS OF AI: POWER, POLITICS, AND THE PLANETARY COSTS OF ARTIFICIAL INTELLIGENCE* 213 (Yale Univ. Press 2021).

<sup>44</sup> John Herrman, *The AI Magic Show*, *INTELLIGENCER* (NEW YORK MAGAZINE), Jan. 18, 2023, <https://nymag.com/intelligencer/2023/01/why-artificial-intelligence-often-feels-like-magic.html> (last visited Jan. 3, 2024) ("OpenAI, which has been accused by its peers of releasing tools to the public with reckless speed, is particularly good at designing interfaces for its models that feel like magic. 'It's a conscious design imperative to produce these moments of shock and awe,' Crawford says. 'We're going to keep having those moments of enchantment.'").

<sup>45</sup> See generally *Mechanical Turk*, WIKIPEDIA, [https://en.wikipedia.org/wiki/Mechanical\\_Turk](https://en.wikipedia.org/wiki/Mechanical_Turk) (last visited Jan. 3, 2024).

combustion engine. The key is understanding the technology *at the appropriate level*. You don't have to be a mechanical engineer to be allowed to drive.

Second, another important reason to understand the basic technological underpinnings of GenAI is that some technologists use various species of a “this is all above you” argument to discredit their critics and to discourage newcomers. Machine learning involves heavy math, and “mathematicians . . . speak of their profession with quasi-religious sentiments and think of themselves as mere prospectors of a transcendental order.”<sup>46</sup> Some have an attitude that lesser minds need not inquire, and should just leave dealing with AI to “the right people.”<sup>47</sup> Do not allow yourself to be intimidated into silence.<sup>48</sup> You do not need to be able to *do the math*; you need only to understand *what the math does*.

With that goal in mind, let's visit the [Restricted Section](#) at Hogwarts or crack open the [Necronomicon](#) and turn the workings of GenAIs from magic and mystery into simplified math and knowledge. GenAI tools today are referred to as models, borrowing that term directly from mathematics. For example, ChatGPT, Bard, Claude, Gemini, and LLaMA are all examples of GenAI models.<sup>49</sup> While all different, these models share some commonalities in their creation and use.

To begin, they have a training phase where typically a *very* large amount of data is run through a computer program which, through some mathematical formulas, learns how the words within that training are statistically connected. The output of this exercise is a

---

<sup>46</sup> Pradeep Niroula, *Bringing the Lofty Ideas of Pure Math Down to Earth*, MIT TECHNOLOGY REVIEW, Apr. 24, 2023, <https://www.technologyreview.com/2023/04/24/1071371/book-reviews-math-education/> (last visited Jan. 3, 2024).

<sup>47</sup> Cf., e.g., Rachel Jang & Daniel Leufer, ‘*Politicians Don't Know Anything About AI, So Governments Can't Regulate It*’, AI MYTHS, <https://www.aimyths.org/we-cant-regulate-ai#politicians-dont-know-anything-about-ai-so-governments-cant-regulate-it> (last visited Jan. 3, 2024) (“The argument usually leads to the conclusion that clueless governments will only make things worse by attempting to regulate AI, so they should leave things to the companies with expertise in the area.”).

<sup>48</sup> Cf., e.g., Matt Laslo & Khari Johnson, *Inside the Senate's Private AI Meeting with Tech's Billionaire Elites*, WIRED, Sept. 14, 2023, <https://www.wired.com/story/senate-ai-forum-musk-zuckerberg/> (last visited Jan. 3, 2024) (“In a historic first . . . upwards of 60 senators sat like school children—not allowed to speak or even raise their hands—in a private briefing where some 20 Silicon Valley CEOs, ethicists, academics, and consumer advocates prophesied about AI's potential to upend, heal, or even erase life as we knew it.”).

<sup>49</sup> GenAI Models encompass a far wider breadth than only the most popular chat models. For a primer on GenAI Models and their other uses, see <https://www.datacamp.com/blog/what-is-a-generative-model> (last visited Jan. 7, 2024).

compressed data set referred to as the “parameters” or “weights” of the model.<sup>50</sup> The objective is for the final model to correctly predict the next word when given some context. The cost and time to complete this training is high, and requires significant computing resources in the form of graphical processing units (“GPUs”). For example, a now relatively small model, LLaMA, took approximately 1 million hours of GPU use, at an estimated cost of \$2.4 million. Much bigger models (meaning they were trained on more information) like ChatGPT-4 have higher training costs.<sup>51</sup>

Once the initial training is completed, a subsequent training can be performed, often referred to as “fine-tuning.” This training follows the same pattern, but instead of a large, unstructured set of data, this second set is highly curated, task-oriented data. For example, for chat models, the fine-tuning data are examples of conversations. Rather than learn specific information, the model instead learns an expected conversation pattern. The fine-tuning process can be repeated and can get very complicated. For our purposes, just know that more fine tuning does not necessarily equate to higher quality responses. It is possible to “over-tune” or “overfit” your model and lower its quality. Essentially, the model becomes too good at predicting the training and loses the ability to generalize things outside of the training data.<sup>52</sup>

This struggle to strike the balance of training and fine-tuned control versus flexibility is not unique to GenAI. Hundreds of years ago, Takuan Sōhō, a Japanese Zen monk and philosopher, espoused the principle he called “seeking the lost mind.” Recognizing the human mind’s tendency to wander and get lost, Sōhō acknowledges that in the early training for the mind, one should seek after and recover this “lost mind.”<sup>53</sup> GenAI models, especially with little training, are also prone to wandering, and their outputs might be interesting but not helpful. However, “[t]he effect of tightening up on the mind is to make it unfree.”<sup>54</sup> The same is true for GenAI models that have been over-trained—

---

<sup>50</sup> See generally <https://deepai.org/machine-learning-glossary-and-terms/weight-artificial-neural-network> (last visited Jan. 7, 2024).

<sup>51</sup> Jonathan Vanian & Kif Leswing, *ChatGPT and Generative AI Are Booming, But the Costs Can Be Extraordinary*, CNET, Mar. 13, 2023, <https://www.cnbc.com/2023/03/13/chatgpt-and-generative-ai-are-booming-but-at-a-very-expensive-price.html> (last visited Jan. 7, 2024).

<sup>52</sup> E.g., Prabhu Srivastava, *Fine-Tuning AI Models: A Guide*, MEDIUM, July 22, 2023, <https://medium.com/@prabhuss73/fine-tuning-ai-models-a-guide-c515bcd4b580> (last visited Jan. 15, 2024).

<sup>53</sup> TAKAUAN SŌHŌ, *THE UNFETTERED MIND* 30 (William Scott Wilson trans., Shambala Publications 2002).

<sup>54</sup> *Id.*

they get stuck on limited data and tasks. Instead, both the Zen master and the GenAI model builder should try to “[b]e in possession of a mind that has been let go of.”<sup>55</sup>

Interestingly and excitingly, once those parameters have been generated and the model creator is satisfied, the computing cost of “inferencing” (using) the model is low. While low costs can become large in the aggregate,<sup>56</sup> for open-source models, private models, and smaller models, this can be done on a single home PC with a graphics card. How do we go from a strange word-association compression of a gargantuan amount of data to the GenAI models we see today, capable of having conversations and rapidly digesting large amounts of text? The key is the “Transformer Neural Network.”

The transformer architecture in a neural network is complicated, but can be understood by recognizing that computers are bad at words, but great at numbers.<sup>57</sup> A Transformer Neural Network takes advantage of a computer’s natural strength. It accepts words as inputs (prompts), converts those words into numbers using the pre-trained and computed information (the parameters), and combines that with positional information numbers (where the word is in the context of a sentence).<sup>58</sup> This produces a resulting number that contains both meaning *and* context. Using this number, and “paying attention”<sup>59</sup> to all of the previously generated words, a Transformer Neural Network predicts the next associated number, then converts the predicted number back into its associated word, adds that new associated word to the output, and repeats that process until it generates a system message representing the end of its sentence or message.

At this point, you can now understand the naming convention behind “ChatGPT”: a **Generative, Pre-trained, Transformer** model, fine-tuned to **chat** with humans. If you are

---

<sup>55</sup> *Id.*

<sup>56</sup> See, e.g., Aaron Mok, *ChatGPT Could Cost Over \$700,000 per Day to Operate. Microsoft is Reportedly Trying to Make it Cheaper*, BUSINESS INSIDER, Apr. 20, 2023, <https://www.businessinsider.com/how-much-chatgpt-costs-openai-to-run-estimate-report-2023-4> (last visited Jan. 15, 2024).

<sup>57</sup> The yin to the lawyer’s yang.

<sup>58</sup> To illustrate just how important positional information and context is in an English sentence, the Finnish Comedian ISMO provides great examples. See ISMO, *I Didn’t Know Sh\*t*, YOUTUBE, <https://www.youtube.com/watch?v=igh9iO5BxBo> (last visited Jan. 13, 2024) (“If I’m a piece of shit, I’m being selfish. And if I’m a giant piece of shit, I’m being more selfish! But if I’m THE shit, I’m great!”).

<sup>59</sup> Vaswani *et al.*, *Attention is All You Need*, *supra* n.32.

still feeling a little lost, we highly recommend the great explanation from the *Financial Times*: “Transformers: the Google scientists who pioneered an AI revolution.”<sup>60</sup>

At an even less technical level, consider the explanation provided by technologist [Jaron Lanier](#) at a recent *Wired* magazine event.<sup>61</sup> His “cartoon model” of how GenAI works uses the classic “is this a cat or a dog?” AI task. We have illustrated his explanations with [Noun Project](#) icons and help from [DALL-E](#).

- Step One: Take a measurement you think will help answer whether something is a cat or a dog. Call each measurement a “neuron.”

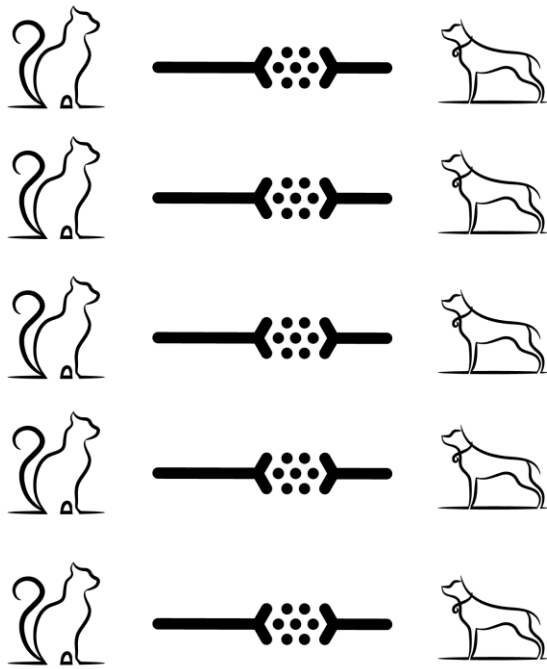


- Step Two: Have multiple layers of “cat-or-dog?” neurons, where each higher layer takes measurements of the lower ones. This is called “deep learning” because the “neurons” are stacked up into towers.

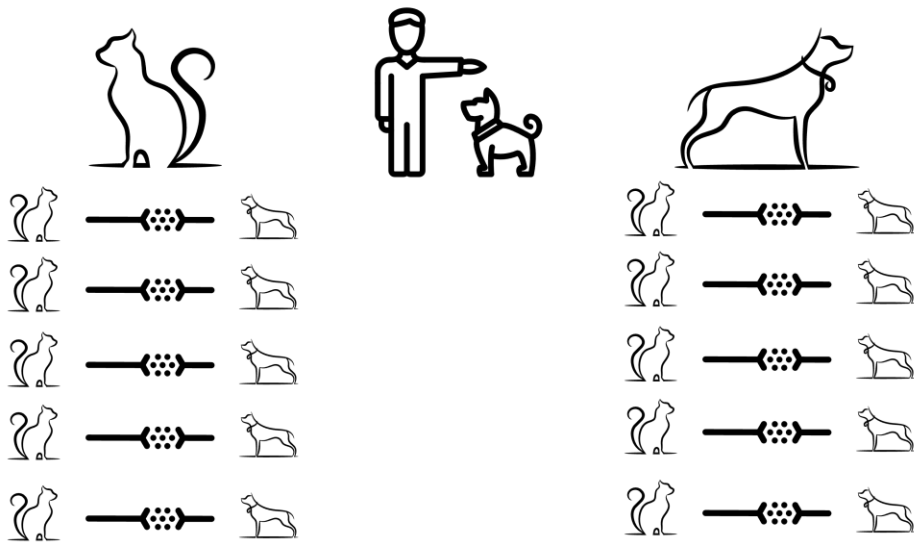
---

<sup>60</sup> Madhumita Murgia, *Transformers: the Google scientists who pioneered an AI revolution*, FINANCIAL TIMES, July 22, 2023.

<sup>61</sup> WIRED, *The Future According to Jaron Lanier*, Dec. 5, 2023, <https://www.wired.com/video/watch/the-future-according-to-jaron-lanier> (last visited Jan. 3, 2024).




- Step Three: “Train” the model by increasing the weights / parameters given to outcomes you deem successful and reducing them for outcomes you deem unsuccessful. After doing it over and over, eventually you get fairly accurate “cat or dog” results from your towers of neurons. This is how a “classifier” AI works.






- Step Four: Repeat for *everything*, so that your model doesn't just have towers of neurons about cats and dogs. It has towers of neurons about *everything*, from parachutes and banjos to cowboy hats and beyond.



- Step Five: Do things in reverse, giving a prompt to the model and asking it to create something new for you out of the trained towers of neurons, rather than classify something existing that you gave to it. E.g., “create an image of a cat.”



- Step Six: Repeat with multiple, diverse prompts, making your model go to multiple towers and blend various outputs together at the same time. E.g., “create

a Van Gogh style lithograph of a cat playing a banjo wearing a cowboy hat falling from the sky with a parachute.” Then you get something like this:



As Lanier sums it up:

So what I propose here is that we can understand what generative AI does . . . as filling in the space between the towers. So what happens is, you have one tower for cats, one tower for parachutes, another tower for banjos. Previously there was a space between them, and now, in order to generate, it fills in that space with something that meets all of [the prompt criteria] at once.

[This “towers” idea] both shows what’s special about [GenAI], but also its limits. It goes up to the level of the towers, but it’s unlikely to go much higher. . . . And so what the technology can do is it can fill in between the

towers, and yet, we shouldn't expect it to grow beyond the towers. . . . All it can do is reconcile towers that we've already defined. . . .

Now that's an incredible value! When you program, you're inevitably, tediously reconstructing something similar to things that people have done millions of times before. And it's painfully tedious. Well, with generative AI, you can leverage all those previous times, and get a version that matches your circumstances and saves you a bunch of time. That's filling in the space between the towers.<sup>62</sup>

If you are with us still, you may already recognize exactly why lawyers are uniquely positioned to take advantage of this powerful new technology. The “towers of neurons” concept shows why GenAI in its current form cannot become super-intelligent and take over the world (it can't go “above the towers”). And it also illuminates what it takes to harness the power of GenAI: substantive knowledge of the meanings of the things in the towers, the verbal abilities to point the tool towards the relevant towers, and foreknowledge of the outcome you are seeking from the tool. It also drives home the point that GenAI operates by trying to statistically predict the outcome you are seeking. It is therefore important to understand what predictive statistics can (and can't) do.

### 3. YOGI BERRA, THE BLACK SWAN, AND THE SORCERER'S APPRENTICE

*“It's hard to make predictions, especially about the future.”*  
– Yogi Berra<sup>63</sup>

With all due respect to Yogi and the future, it turns out that it's also hard to make predictions about *anything* that really matters. GenAI (like most forms of AI) is inherently a probabilistic, predictive technology.<sup>64</sup> These methods are often suggested as

---

<sup>62</sup> *Id.*

<sup>63</sup> The quote is attributed to many people, including but not limited to Berra. *See, e.g.*, <https://quoteinvestigator.com/2013/10/20/no-predict/> (last visited Jan. 3, 2024).

<sup>64</sup> *The Future According to Jaron Lanier, supra* n.61 (“Another friend of mine wrote a piece . . . talking about how all AI can do is regurgitate randomly. Timnit [Gebu] called it ‘stochastic parrots’ . . . . Anyway, so this is true. This is what it does.”).

ways to deal with uncertainty.<sup>65</sup> That uncertainty could be anything from the probability that your fantasy football team wins its matchup this week to the probability that a client wins litigation. As discussed above regarding transformers, it could also be the probability of the next letter, next word, or next sentence that is responsive to a prompt.

When faced with uncertainty, it is intuitive that existing data could be probative of unknown data, and the past probative of the future.<sup>66</sup> In some sense, one can attempt to reduce anything to a prediction based on probabilities.<sup>67</sup> Consider the existence of planet Earth. We have approximately 4.3 billion years of data on planet Earth. Each day in that data set, the sun has risen in the morning. We can probably feel confident in predicting, with “five nines”<sup>68</sup> level of confidence (that is, 99.999% certain), that tomorrow the sun will rise. It is therefore tempting to resolve uncertainty by making decisions based on the previously observed frequency of an event. After all, if past evidence shows that a thing is very unlikely to happen, why worry about it?

Alas, this is a major logical mistake. And so is concluding that a thing which happens all the time is no big deal. This is because for purposes of making decisions under conditions of uncertainty, *the consequences of an event are far more important than its frequency.*<sup>69</sup>

This is true in two very different ways. First, let’s consider the really big consequences that rarely occur. Noted scholar of uncertainty [Nassim Taleb](#) calls these “Black Swans.”<sup>70</sup> For example, based on all the scientific evidence we have at present, we can say that

---

<sup>65</sup> PETER L. BERNSTEIN, *AGAINST THE GODS: THE REMARKABLE STORY OF RISK* 4–5 (John Wiley & Sons 1998) (noting “the Law of Large numbers and methods of statistical sampling that drive modern activities as varied as opinion polling, wine tasting, stock picking, and the testing of new drugs”).

<sup>66</sup> *See, e.g., id.* at 121 (noting Jacob Bernoulli’s insights, driven by the “giant assumption” that “under similar conditions, the occurrence (or non-occurrence) of an event in the future will follow the same pattern as was observed in the past” (quoting Bernoulli himself)).

<sup>67</sup> Indeed, as far as we know from the current state of quantum mechanics, at the atomic-particle level everything simply *is* just a prediction.

<sup>68</sup> *Five Nines*, PC MAG, <https://www.pcmag.com/encyclopedia/term/five-nines> (last visited Jan. 3, 2024).

<sup>69</sup> NASSIM NICHOLAS TALEB, *THE BLACK SWAN: THE IMPACT OF THE HIGHLY IMPROBABLE* 211 (Random House Trade paperback ed. 2021) (“This idea that in order to make a decision you need to focus on the consequences . . . rather than the probability . . . is the *central idea of uncertainty.*” (emphasis in original)).

<sup>70</sup> “[A] Black Swan . . . is an event with the following three attributes[:] First, it is an *outlier* . . . . Second, it carries an extreme impact . . . . Third, in spite of its outlier status, human nature makes us concoct explanations for its occurrence *after* the fact . . . .” *Id.* at xxii (emphasis in original).

someday, the sun will not rise.<sup>71</sup> Similarly, the possibility of an asteroid strike is vanishingly small on a day-to-day basis.<sup>72</sup> We can probably say with “five nines” level of confidence that no asteroid strike will occur tomorrow. But the consequences of such an event vastly outweigh its low likelihood. Just ask the dinosaurs.

Second, there are also individually miniscule consequences that, in the aggregate, add up to big impacts. The old adage “death by a thousand cuts” expresses this well. For example, a loss of a single dollar is irrelevant. A loss of a single dollar on each of a million transactions is a shellacking.

One may well ask why the consequences of an event cannot also be statistically predicted, in the same manner as the frequency of an event. The answer is that statistical predictions assume that we have an unbiased data set to which those statistical methods can be applied—a challenge that concerns both frequency and consequences.<sup>73</sup> In many real-world situations, and particularly with regard to consequences, “past data from real life constitute a sequence of events rather than a set of independent observations, which is what the laws of probability demand.”<sup>74</sup> We have no way of *really* knowing whether our sample is adequate to support our probabilistic predictions.<sup>75</sup>

GenAI is very good at predicting the next word, especially when the prompt or input is close to the model’s training data. Unfortunately, the level of prediction is nowhere close to “five nines” of certainty.<sup>76</sup> While it is possible to adjust the level of variability in most

---

<sup>71</sup> Whether this occurs by the sun becoming a red giant that eats the earth, some other astronomical cataclysm, or simply the heat death of the entire universe, all evidence points to this eventual inevitability.

<sup>72</sup> E.g., Alan Harris, *The Odds of an Asteroid Strike*, NOVA, Mar. 27, 2023, <https://www.pbs.org/wgbh/nova/article/risk-of-an-asteroid-strike/> (last visited Jan. 3, 2024).

<sup>73</sup> E.g., TALEB, *THE BLACK SWAN*, *supra* n.69, at 281 (“[Y]ou need only one single observation to reject the Gaussian, but millions of observations will not fully confirm the validity of its application.”); BERNSTEIN, *AGAINST THE GODS*, *supra* n.65, at 334 (“Thus, forecasting tools based on . . . computer gymnastics are subject to many of the same hurdles that stand in the way of conventional probability theory: the raw material of the model is the data of the past.”).

<sup>74</sup> BERNSTEIN, *AGAINST THE GODS*, *supra* n.65, at 335.

<sup>75</sup> E.g., TALEB, *THE BLACK SWAN*, *supra* n.69, at 41 (“[This problem] strikes at the nature of empirical knowledge itself. Something has worked in the past, until—well, it unexpectedly no longer does, and what we have learned from the past turns out to be at best irrelevant or false, [and] at worst viciously misleading.”).

<sup>76</sup> Accuracy of word prediction is nearly impossible to measure for GenAI, because even answering the question “what was correct” is mostly subjective. For a discussion of which models are the most accurate in

models' predictions,<sup>77</sup> even the “most certain” level of variability is still prone to incorrect or fabricated answers. And even more problematic, each GenAI variance or deviation is compounded: every predicted word is *fed back into the next predicted word*, so a single incorrect or unintended prediction can derail the entire response. It is very similar to the children's game of telephone, but with more serious consequences.

Some will understandably be tempted to take a leap of faith here. AI maximalists will note all of the benefits that can come from turning over tasks to predictive machines. Especially for businesses (and each lawyer, even in a firm, is really their own small business), profitability can in theory be greatly increased by outsourcing as many tasks as possible to AI. “Sure, AI does not always work as desired,” they may say, “but neither do humans. And you have to weigh the risk of adverse consequences versus the reward of greater profit margins from efficiency gains. No risk, no reward.”

While true, this assumes that the consequences can be appropriately quantified and weighed against the possible rewards. Human relationships are nuanced and multifaceted, and human-relationship consequences are not readily reducible to numbers. For example, an early use of predictive technologies was autocorrect—the foibles of which gave rise to an entire book (*[Damn You, Autocorrect!](#)*). What exactly are the dollar value consequences of an AI accidentally telling your client something like this:

---

question-answers, see <https://originality.ai/blog/what-llm-is-the-most-accurate> (last visited Jan. 15, 2024). Additionally, GenAI models, during their creation, can be configured on what words should be predicted through training data and math. For an in-depth explanation of this process using only Shakespeare as the training data, see Andrej Karpathy, *Let's Build GPT: from scratch, in code, spelled out.*, YOUTUBE, <https://www.youtube.com/watch?v=kCc8FmEb1nY> (last visited Jan. 14, 2024).

<sup>77</sup> Most GenAI models have an input known as ‘temperature’ which allows for the user to dial up or down variability of the response. For more consistent responses, a low temperature is used to force the model to always select the words with the highest likelihood, while for creative responses, a high temperature is used for less likely (and therefore, more random) words to be selected. See David Eastman, *What Temperature Means in Natural Language Processing and AI*, THE NEW STACK, Jan. 9, 2024, <https://thenewstack.io/what-temperature-means-in-natural-language-processing-and-ai/> (last visited Jan. 15, 2024).



All of these things point to a simple but fundamental limitation on all current forms of artificial intelligence, GenAI included: while their statistical methods may be *competent* at certain tasks, they lack *comprehension* of the human meanings of those tasks.<sup>79</sup> This distinction has been known for quite some time with “traditional” forms of classifying AI,<sup>80</sup> and it holds true for GenAI as well.

Indeed, the tasks at which traditional AI is competent can be quite similar to those where human beings “shoot from the hip” without deliberation. This is what Nobel winner Daniel Kahneman calls “System 1” thinking, which is largely unconscious, heuristic-based, and thus prone to biases and errors.<sup>81</sup> Consider again the “cat or dog” classifying

---

<sup>78</sup> Jillian Madison, ‘*Damn You Autocorrect*’: *Funniest iPhone Autocorrect Mistakes*, HUFFINGTON POST, Mar. 22, 2011, [https://www.huffpost.com/entry/funniest-iphone-autocorrect-mistakes\\_b\\_838684](https://www.huffpost.com/entry/funniest-iphone-autocorrect-mistakes_b_838684) (last visited Jan. 3, 2024).

<sup>79</sup> The competence / comprehension distinction is commonly attributed to philosopher Daniel Dennett, who in turn drew upon legendary computer scientist Alan Turing. Dan Jones, *Cognitive Science: Dennett Rides again*, NATURE, Feb. 2, 2017, <https://www.nature.com/articles/542030a> (last visited Jan. 3, 2024) (“Dennett also draws heavily on the idea of ‘competence without comprehension,’ best illustrated by mathematician Alan Turing’s proof that a mechanical device could do anything computational.”).

<sup>80</sup> E.g., Stephen M. Fleming, *What Separates Humans from AI? It’s Doubt*, FINANCIAL TIMES, Apr. 15, 2021 (“The real danger . . . [is] that we will *over-estimate* the comprehension of our latest thinking tools, prematurely ceding authority to them far beyond their competence.” (quoting Jack Maden, *What Happens When Machines Become Smarter Than People?*, PHILOSOPHY BREAK, Dec. 2019)).

<sup>81</sup> DANIEL KAHNEMAN, *THINKING, FAST AND SLOW* 24–25 (Farrar, Straus & Giroux paperback ed. 2013) (“System 1 is generally very good at what it does: its models of familiar situations are accurate, its short-term predictions are usually accurate as well, and its initial reactions to challenges are swift and generally appropriate. System 1 has biases, however, systematic errors that it is prone to make in specified circumstances.”).

task. AI has become competent at this.<sup>82</sup> Human beings are highly competent at it, without ruminating over the question. You “just know” that something is a cat, and almost instantaneously.

But in questionable cases where it really matters, human beings can call upon higher-level “System 2” thinking: slow, deliberate, reasoned, conscious, and intensive.<sup>83</sup> Current AIs lack the ability to engage in System 2 thinking—or even to determine whether or not it is worth doing so. *Comprehension* of the task’s meaning and importance is the hallmark of System 2. As one scholar notes:

I believe the main remaining gap between current advanced systems and [“true” artificial general intelligence] is what we could refer to as conscious cognition—abilities such as reasoning, deliberate thought and explicit planning. I have argued for many years that although deep learning has made huge strides in cognitive capabilities corresponding to human intuition (system 1), methods are still weak regarding the conscious cognition crucial for humans to provide correct answers in settings for which we need to reason (system 2).<sup>84</sup>

For human activities that really matter (like legal work), judging the *consequences* of a thing requires having *comprehension* of the thing. High-consequence issues need more than shoot-from-the-hip statistical shortcuts. As a famous paper put it, GenAI can be thought of as a “stochastic parrot”—it gives human-sounding answers in human-sounding language, but has no idea what any of it *means*:

Contrary to how it may seem when we observe its output, [a large language model] is a system for haphazardly stitching together sequences of linguistic forms it has observed in its vast training data, according to

---

<sup>82</sup> E.g., Aditya Srinivas Menon, *Classifying Cats vs Dogs: A Beginner’s Guide to Deep Learning*, MEDIUM, Nov. 28, 2020, <https://towardsdatascience.com/classifying-cats-vs-dogs-a-beginners-guide-to-deep-learning-4469ffed086c> (last visited Jan. 3, 2024).

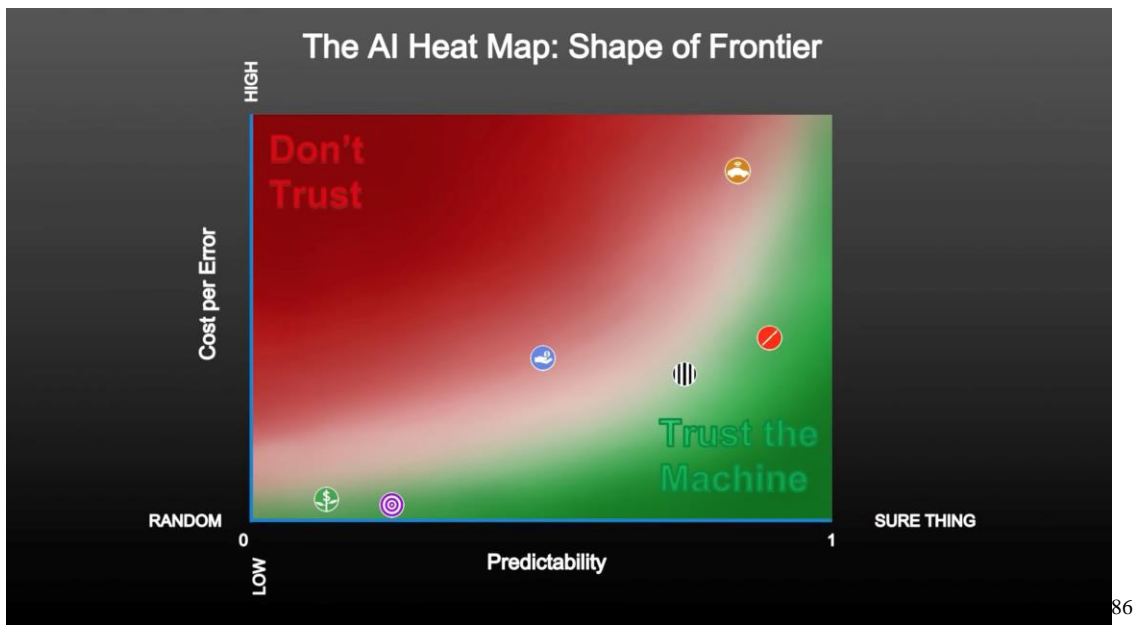
<sup>83</sup> KAHNEMAN, THINKING, FAST AND SLOW, *supra* n.81, at 21 (“System 2 [is] the conscious, reasoning self that has beliefs, makes choices, and decides what to think about and what to do.”); *id.* at 24 (“When System 1 runs into difficulty, it calls on System 2 to support more detailed and specific processing that may solve the problem of the moment . . . . System 2 is activated when an event is detected that violates the model of the world that System 1 maintains.”).

<sup>84</sup> Yoshua Bengio, *For True AI Governance, We Need to Avoid a Single Point of Failure*, FINANCIAL TIMES, Dec. 4, 2023 (opinion of a Turing Award winning scientist).



probabilistic information about how they combine, but without any reference to meaning: a stochastic parrot.<sup>85</sup>

So while GenAI is *very* competent at parroting out variations of words, it has no comprehension of the consequences that can arise from those words. This competence / comprehension divide creates a risk / reward curve for delegating any task to an AI, including a GenAI. [Vasant Dhar](#), a noted AI scholar, has explained that this challenge can be plotted on a graph having two elements: the statistical predictability of the task, and the consequences of making an error in the task:



This brings us to the cardinal rule of using GenAI in legal work:

**Consequences control: in anything that truly matters,  
AI must never be left unsupervised.**

<sup>85</sup> See Emily M. Bender *et al.*, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?* 2021 PROCEEDINGS OF THE ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY (Mar. 2021), at 616–17, available at <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922> (last visited Jan. 3, 2024).

<sup>86</sup> Vasant Dhar, *When Should We Trust Machines?*, TED, Feb. 2018, [https://www.ted.com/talks/vasant\\_dhar\\_when\\_should\\_we\\_trust\\_machines](https://www.ted.com/talks/vasant_dhar_when_should_we_trust_machines) (last visited Jan. 3, 2024).

No matter how tempting or cost-effective, a task with any material importance cannot be turned over to any AI (GenAI included) without a competent lawyer reviewing the AI's outputs. This is true from both a legal-ethics standpoint and a purely practical standpoint.

On the ethics point, as we write this paper, the bench and bar are only beginning to address this topic, but the early results are in accord. For example, the California state bar's professional responsibility committee recently stated:

A lawyer's professional judgment cannot be delegated to generative AI and [that judgment] remains the lawyer's responsibility at all times. A lawyer should take steps to avoid over-reliance on generative AI to such a degree that it hinders critical attorney analysis . . . .<sup>87</sup>

By the time you are reading this paper, we predict that more ethics authorities will have coalesced around this rule.

But wholly apart from the legal-ethics considerations, basic client service demands it. For justifiable economic reasons, law firms have long “pushed down” work to the lowest-rate biller. But that biller must know what they are doing: no matter how cost-effective it may be to completely hand off a complex transaction to a first-year paralegal, doing so courts malpractice catastrophe and wagers the loss of the client—unacceptable consequences.<sup>88</sup> And although GenAI writes convincingly and can pass the multistate bar exam,<sup>89</sup> when unsupervised it is not even to the level of that first-year paralegal, because GenAI does not understand the social meanings of any of the tasks it performs or the words it writes.

---

<sup>87</sup> *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law (Duties of Competence and Diligence)*, STATE BAR OF CALIFORNIA, STANDING COMMITTEE ON PROFESSIONAL RESPONSIBILITY AND CONDUCT, available at <https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf> (last visited Jan. 13, 2024).

<sup>88</sup> *Cf.* MODEL RULE OF PROFESSIONAL CONDUCT 5.3 cmt. (“A lawyer must give such assistants appropriate instruction and supervision . . . and should be responsible for their work product.”).

<sup>89</sup> Pablo Arrendondo *et al.*, *GPT-4 Passes the Bar Exam: What That Means for Artificial Intelligence Tools in the Legal Profession*, STANFORD CENTER FOR LEGAL INFORMATICS, Apr. 19, 2023, <https://law.stanford.edu/2023/04/19/gpt-4-passes-the-bar-exam-what-that-means-for-artificial-intelligence-tools-in-the-legal-industry/> (last visited Jan. 3, 2024).

Of course, while we feel quite confident in holding that nothing of **material consequence** should be turned over to an **unsupervised AI**,<sup>90</sup> what is “material” and what “supervision” entails varies considerably depending on the facts. The exact same task may be material to one, but not to another. Consider a boring, laborious, burdensome task where the known, maximum possible adverse consequences of a mistake are \$10,000.<sup>91</sup> That amount is certainly material to us, and we would not outsource the task to any currently existing GenAI. For Warren Buffett, the result is probably the opposite.

But even for Mr. Buffett, a task that is individually immaterial may become material when done “at scale.” As an old joke about human beings goes, “Cut me a break, I’m making mistakes as fast as I can.” Machines are not so constrained, and can process thousands of decisions per second. If the machine is making mistakes in doing so, what was a minimal consequence individually can quickly become material in the aggregate. A \$10,000 mistake repeated millions of times becomes material almost all the time.<sup>92</sup>

We call this the “Sorcerer’s Apprentice problem,” after the famous segment in Disney’s classic *Fantasia*. In it, Mickey Mouse casts a spell that turns a broom into a magical worker, freeing him from the boring, laborious, burdensome task of fetching water from a well. Immensely self-satisfied, Mickey takes a nap only to discover that the magical broom—lacking any comprehension of what it was doing—has flooded the whole place by continuing to bring water unendingly.<sup>93</sup>

---

<sup>90</sup> See, e.g., Michael O’Dwyer et al., *PwC Experiments with Chatbots to Speed Up Its Lawyers*, FINANCIAL TIMES, Mar. 14, 2023 (“The firm said the technology would speed up decision-making by producing answers to questions, which would then be reviewed and added to by staff.”).

<sup>91</sup> It is worth asking how often one can really know, with certainty, the maximum consequences of an adverse outcome—especially in precise dollar amounts. The challenge of pinning down precise consequences adds to the uncertainty of making the decision to outsource work. But here the hypothetical’s job is simply to make the point.

<sup>92</sup> This why Nassim Taleb states: “Principle: A unit needs to take any risk as if it were going to take it repeatedly—at a specified frequency—over its remaining lifespan.” NASSIM NICHOLAS TALEB, *SKIN IN THE GAME: HIDDEN ASYMMETRIES IN DAILY LIFE* 249 (Random House Trade Paperback ed. 2021) (emphasis removed). And further underscoring the importance of *consequences rather than frequency* as the key criterion for risk-based decision-making, eliminating the consequence of ruin is key: “While experiments are static . . . life is continuous. If you incur a tiny probability of ruin as a ‘one-off’ risk, survive it, then do it again (another ‘one-off’ deal), you will eventually go bust with probability 1. Confusion arises because it may seem that the ‘one-off’ risk is reasonable, but that also means that an additional one is reasonable.” *Id.* at 249–50.

<sup>93</sup> SORCERER’S APPRENTICE – FANTASIA (Disney 1940), <https://video.disney.com/watch/sorcerer-s-apprentice-fantasia-4ea9ebc01a74ea59a5867853> (last visited Jan. 3, 2024).

In short, “AI may be tactically brilliant but it is strategically banal.”<sup>94</sup> Strategy remains the realm of human beings who comprehend the social meanings of the tasks being performed. The human leveraging a GenAI now has much in common with a symphony conductor or a film director,<sup>95</sup> whose job it is to holistically shape the actions of specialized and finicky performers. This has crucial implications for understanding why seasoned lawyers are well suited to wield the technology.

#### 4. PROGRAMMING IN NATURAL LANGUAGE

*“Just because something doesn’t do what you planned it to do doesn’t mean it’s useless.”*

– Thomas Edison <sup>96</sup>

So far we have established the following propositions about GenAI:

- It is not human;
- It runs on math, not magic;
- It may competently do something, but it lacks comprehension of what it does.

Now we will establish a fourth important proposition: Because GenAI does not do anything without “prompting,” the ability to prompt it in ways that generate a response that you desire is critical. Prompting points GenAI to the relevant “towers of neurons” for it to fill in. The ability to carefully craft a prompt to get an intended outcome is called “prompt engineering.” And since prompt engineering is done in natural language, the

---


<sup>94</sup> LAWRENCE FREEDMAN, *COMMAND: THE POLITICS OF MILITARY OPERATIONS FROM KOREA TO UKRAINE* 506 (2022).

<sup>95</sup> See, e.g., Kevin Kelly, *Picture Limitless Creativity at Your Fingertips*, WIRED, Nov. 17, 2022, <https://www.wired.com/story/picture-limitless-creativity-ai-image-generators/> (last visited Jan. 3, 2024) (“It seems obvious to me that promptors are making true art. What is a consummate movie director—like Hitchcock, like Kurosawa—but a promptor of actors, actions, scenes, ideas? Good image-generator promptors are engaged in a similar craft . . .”); Will Douglas Heaven, *Generative AI Is Changing Everything. But What’s Left When the Hype Is Gone?*, MIT TECHNOLOGY REVIEW, Dec. 16, 2022 <https://www.technologyreview.com/2022/12/16/1065005/generative-ai-revolution-art/> (last visited Jan. 3, 2024) (“‘The creativity we see today comes from the use of the systems, rather than from the systems themselves,’ says Llano—from the back-and-forth, call-and-response required to produce the result you want.”).

<sup>96</sup> E.g., [https://www.brainyquote.com/quotes/thomas\\_a\\_edison\\_100430](https://www.brainyquote.com/quotes/thomas_a_edison_100430) (last visited Jan. 3, 2024).

ability to wield written language with conceptual clarity and precision is vital to making use of GenAI. As it happens, these are precisely the skills of lawyers.

As discussed above, the key innovation behind GenAI is a merger of the “deep learning” branch of neural networks with a natural language interface.<sup>97</sup> The natural language interface means that powers previously available only to hard-core coders are now at the fingertips of nearly anyone.<sup>98</sup> The CEO of the (currently) leading GenAI firm touts this:



We know that many people who want to build a GPT don’t know how to code. We’ve made it so you can program a GPT just by having a conversation. We believe that natural language is going to be a big part of how people use computers in the future . . . .<sup>99</sup>

---

<sup>97</sup> Kelly, *Picture Limitless Creativity*, *supra* n.95 (“AI image generators were born from the marriage of two separate technologies. One was a historical line of deep learning neural nets that could generate coherent realistic images, and the other was a natural language model that could serve as an interface to the image engine. The two were combined into a language-driven image generator.”).

<sup>98</sup> Heaven, *Generative AI Is Changing Everything*, *supra* n.95 (“‘It’s uplifting for all the folks who’ve never been able to create because it was too expensive or too technical,’ he says. ‘But it’s terrifying if you’re not open to change.’”).

<sup>99</sup> OPENAI, *Dev Day 2023: Opening Keynote*, <https://youtu.be/U9mJuUkhUzk?t=1561> (last visited Jan. 3, 2024).

Other sources are in accord.<sup>100</sup> And all of this leads us to our core point:

**Words matter: GenAI is best wielded by those with the ability to write natural language—not code—with deeply reasoned precision.**

Because GenAI lacks “System 2” processing capabilities, leveraging it safely means *the user must supply the System 2 reasoning capabilities that the GenAI lacks*. In our society, the people who know how best to wield words with precision and a view towards the consequences that flow from words are . . . lawyers! Therefore, far from eliminating lawyers, GenAI stands to make the skills of lawyers more important—not less.

To bring this point home, consider a bar trivia quiz as an analogy for using GenAI. The trivia host asks the questions to all the teams in the room, and each of those teams independently answers using their own knowledge without consulting the other teams. In GenAI, the user is the host, asking questions, which the GenAI “inferences” from what it knows (remember the towers of neurons) and produces an answer. The variability (uncertainty) in the answer produced by the GenAI is analogous to the variability or uncertainty you would get by sampling answers from various trivia teams.

For example, say you start with an easy, warm-up round. Most trivia teams are likely to answer easy questions correctly off the top of their heads. If you had to randomly select from ten trivia teams’ answers, your chances of that randomly selected answer being the correct one might be, say, 80–90%. Easy, “off the top of your head” questions are the equivalent of System 1 reasoning questions for GenAI. If you ask System 1 level questions to GenAI, you might get the right answer.

But what if you grab a pint and step up the difficulty of the questions in the bar trivia? You are bound to see a wider variety of answers from the teams: some correct, some wrong, some funny, or some mix of all the above. A randomly selected answer from those ten teams is not likely to be correct. Similarly, if you ask the equivalent of a hard question for GenAI—a System 2 reasoning question—you might receive the correct answer, but you may also receive something made up or funny.

---

<sup>100</sup> Joanna Stern, *Talking to Chatbots is Now a \$200k Job. So I Applied.*, WALL ST. JOURNAL, Nov. 29, 2023 (quoting a startup CEO as saying, “You can think of prompt engineering as programming in the English language.”).

Consider one final step that would never happen at bar trivia. How well do you think the ten trivia teams would do on the same hard question with guidance or *prompting*? Imagine that you could explain to each team the context of the question, handing them background materials to help reference the right answer, and that you gave them step-by-step instructions on how to answer the question. Each team would be armed with the right knowledge and context. They would no longer be answering “off the top of their heads” and **would be able to answer significantly more difficult questions more accurately**. *That* is the opportunity and power that legal thinking and knowledge can bring to GenAI. We lawyers know how to provide the facts, the theories or rules that apply, the potentially distinguishing factors, *and* our question(s), all together. We constantly sharpen these skills against other lawyers, and when used in a prompt for GenAI, these skills can achieve impressive results.

The potential applications are nearly limitless, but clarity of verbal reasoning is key. And lawyering is indivisible from the high-grade, “System 2,” deeply-reasoned processing of words. Indeed, the more that GenAI infiltrates daily life, the more valuable that legal training may become, even for those who never intend to practice law. Thinking like a lawyer means choosing “the right words,” and carefully.

What *are* the right words?<sup>101</sup> Quality, quantity, and precision of words have a dramatic effect on the efficacy of GenAI. As real estate attorneys, we are no strangers to “magic words,” but unlike in our domain, the magic words are not yet readily known for GenAI. Much is yet to be written and discovered in this area. Prompt engineering is an emerging discipline, and it is too soon to provide a comprehensive how-to guide or say what is commonly-accepted practice. GenAI can respond like an overeager summer associate, but that also provides a great rule of thumb: give the same instructions to GenAI as you would to a helpful but naïve summer associate. Here are some basic principles of prompt engineering that seem to be effective and well-regarded as of this writing:

- Use a “Persona Prompt.” As noted above, leveraging GenAI is not unduly different from being the director of a film. A film director must provide a well-ordered setting for actors to do their work. By telling the GenAI to assume a role (as if an actor on a stage), and providing context about the scene, the model is

---

<sup>101</sup> As we emerge from the Restricted Section, we may hear Hermione remind us to use “LeviOsa, not LeviosA,” or as we try to speak the words to open the Necronomicon, we had better remember what comes after “Klaatu Barada” and not mumble the final word as Ash did.

better able to find the correct “towers” where the gaps are to be filled.<sup>102</sup> The more specialized the director, the better the results. Imagine you are selecting from writings across the internet: if you are doing basic legal work, you probably want to mirror things written by other lawyers. “You are a bankruptcy lawyer...” queues GenAI to respond in that style. Analyzing a summary judgement motion you have written? “You are a judge with expertise in real estate, assigned to decide a motion . . . .”

- Use “New or Specific Information Prompts.” As set forth above, GenAIs are trained on data (typically from the internet). However, that data is time-bound or may be without context. A human being continues to gather and sort knowledge as the arrow of time advances, but a GenAI’s knowledge stops when its training data set is established.<sup>103</sup> Although there are ways of fine-tuning a GenAI with additional data sets relevant to a particular use case or context, these too are time-bound. Therefore, sometimes the savvy prompt engineer must provide the GenAI with new information that is not in its training data.<sup>104</sup> Similarly, consider a model that has consumed the entire internet’s information on laws from all fifty states throughout all history. Without direction to the laws that apply to your specific question, the chances that a GenAI selects “the right law” are quite low. Avoid this by providing directed, detailed information into the prompt. Don’t be shy about copy-pasting a statute or a case directly into the prompt if appropriate.
- Use a “Question-Refinement Prompt.” Ask the GenAI to help you sharpen your own question. With “the compressed internet” in their memory banks, GenAIs have access to an astounding number of words. This makes them capable drafting companions that can help you craft or refine a turn of phrase, even one meant to be asked back to the GenAI itself. Consider: “When I ask a question about [topic X], suggest a better version of my question, and then ask me if I would like to use your question instead.”<sup>105</sup>

---

<sup>102</sup> Stern, *Talking to Chatbots*, *supra* n.100.

<sup>103</sup> *E.g.*, OpenAI, Dev Day 2023, *supra* n.99 (“We are just as annoyed as all of you, probably more, that GPT-4’s knowledge about the world ended in 2021. We will try to never let it get that out of date again.”).

<sup>104</sup> Stern, *Talking to Chatbots*, *supra* n.100.

<sup>105</sup> *Id.*



- Ask specific, leading questions. Consider GenAI to be a semi-hostile witness: without careful direction, you never know what it may say.<sup>106</sup> In general, a GenAI *will* attempt to answer any question it gets, even if that means completely fabricating an answer (“hallucinations”). The more specific and leading your question, the more useful the response you may get.
- But not *too* specific or *too* leading.... As a semi-hostile witness without comprehension of the topics on which it is responding, a GenAI will respond only to the letter of your question. It can’t understand the spirit of your question.<sup>107</sup> You have to give it the right amount of leash or it can’t use its probabilistic capabilities to help you. As Sōhō reminds us, an overly-restricted mind is unfree.
- If at first you don’t succeed, try again. GenAIs are probabilistic models, and so the same prompt can sometimes generate different answers. (Frustrating, but true. Remember, it’s designed that way.) And slight changes in wording can make a big difference.<sup>108</sup>
- Break complex things down into simple things. Because GenAI has no System 2 thinking, you **must** give explicit instructions for how to work through detailed chains of reasoning. Don’t leave anything to assumption.<sup>109</sup> (Which is a lot like the process of good legal draftwork . . . .)

Ultimately, the practice of prompt engineering is the art of coaxing the GenAI stochastic parrot to return a result that is closer and closer to your desired one. Skilled prompt engineers can thus “load the dice” of the probabilistic algorithms underlying the GenAI, tilting the odds in their favor. The more effective the prompt engineer, the more efficient the human-machine collaboration between user and large language model.

---

<sup>106</sup> Cf. Irving Younger, *The 10 Commandments of Cross Examination*, <https://www.youtube.com/watch?app=desktop&v=dBP2if0l-a8> (last visited Jan. 2, 2024) (“Always ask leading questions.”); *id.* at 8:45 (“On cross-examination, you will *never, never, NEVER* ask anything *but* a leading question! You put words in the witness’s mouth!!! That’s the whole idea on cross examination!”).

<sup>107</sup> Cf. *id.* (“Don’t ask the ‘one question too many.’”); *id.* at 37:00 (“Well how come!?! If your back was to them . . . how come you know . . . that the defendant, my client, bit off the victim’s nose!?! . . . And the famous answer is, ‘Well, I saw him spit it out.’”).

<sup>108</sup> See, e.g., Jamie Teevan, *To Work Well with GenAI, You Need to Learn How to Talk to It*, HARVARD BUSINESS REVIEW, Dec. 15, 2023, <https://hbr.org/2023/12/to-work-well-with-genai-you-need-to-learn-how-to-talk-to-it> (last visited Jan. 2, 2024).

<sup>109</sup> See, e.g., *id.*

## 5. WHAT LIES AHEAD

*“Come writers and critics, who prophesize with your pen,  
And keep your eyes wide, the chance won’t come again,  
But don’t speak too soon, for the wheel’s still in spin,  
And there’s no tellin’ who that it’s namin’...”*

*– Bob Dylan* <sup>110</sup>

In the wake of GenAI’s explosion onto the scene, it has become fashionable to predict the end of many jobs,<sup>111</sup> including those of lawyers.<sup>112</sup> Some have even predicted the end of the world.<sup>113</sup> So prevalent is the idea that AI could be a powerfully destructive force that the concept has been dubbed “AI doomerism.”<sup>114</sup>

At this stage of the hype cycle, it is difficult to predict granular outcomes with any real confidence. That being said, there are discernable trend lines. These arise from the inherent nature of GenAI discussed above, together with larger structural economic forces, and these provide a more reliable foundation for analysis.

First, the production of first drafts will move away from junior-lawyer humans and towards GenAI models. First-drafting work is time-consuming and low-risk, and thus ripe for disruption by the textual capabilities of GenAI. Further, much first drafting work is simply mimicry of prior trusted documents, which again plays straight into the best capabilities of GenAI. With well-trained GenAI tools at their disposal, a senior lawyer with sharp prompt engineering skills will be able to generate in minutes what previously would have required hours of junior lawyer labor. Today, it is entirely possible for skilled

---

<sup>110</sup> THE TIMES THEY ARE A-CHANGIN’ (Columbia Recording 1964).

<sup>111</sup> E.g., Angela Watercutter, *The Hollywood Strikes Stopped AI From Taking Your Job. But for How Long?*, WIRED, Dec. 25, 2023, <https://www.wired.com/story/hollywood-saved-your-job-from-ai-2023-will-it-last/> (last visited Jan. 3, 2024) (“Throughout 2023, many trades and professions, from painters to coders and beyond, found themselves vulnerable to being replaced by machine learning.”).

<sup>112</sup> E.g., <https://willrobotstakemyjob.com/lawyers>.

<sup>113</sup> E.g., <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>.

<sup>114</sup> E.g., Melissa Heikkila, *Four Trends that Changed AI in 2023*, MIT TECHNOLOGY REVIEW, Dec. 19, 2023, <https://www.technologyreview.com/2023/12/19/1085696/four-trends-that-changed-ai-in-2023/> (last visited Jan. 3, 2024) (calling 2023 the year that “AI doomerism went mainstream”).

senior lawyers to create “MeGPT” versions of themselves, trained on their own preferred documents and writing styles, that can do their drafting bidding on command.<sup>115</sup>

This may have significant implications for the work lives of in-training or early-career lawyers. For decades, new lawyers and large law firms have engaged in a symbiotic trade. Roughly speaking, in exchange for working long hours on short notice doing “grunt work” like first drafts, the early-career lawyer would acquire substantive knowledge and practice area expertise. “Learning by osmosis” over a 3–5 year job tenure of reading and attempting to create substantive documents that are reviewed and edited by senior lawyers, junior lawyers would increase their skills. After time in the trenches, the upskilled associate could then continue building a practice at the large firm, or take these new skills to other employers (smaller firms, solo practice, in-house jobs, etc.).<sup>116</sup>

While this piece of the large-firm career model has been unstable and criticized for a long time,<sup>117</sup> GenAI portends to upend it. Simply put, large law firms are for-profit businesses, and did not set out to be eleemosynary institutions. In a world where GenAI can create first drafts of even complex documents for pennies in compute costs, and where clients are dissatisfied with being charged time to train associates,<sup>118</sup> it may well undercut the economic reason to farm first-draft work out to junior lawyers, and to mentor and upskill those junior lawyers. Faced with these economics, it will take “something more” to be an attractive candidate to a large law firm in the near future.<sup>119</sup>

---

<sup>115</sup> See, e.g., Reece Rogers, *How to use OpenAI’s ChatGPT to Create Your Own Custom GPT*, WIRED, Dec. 26, 2023, <https://www.wired.com/story/how-to-use-chatgpt-create-custom-gpt-openai/> (last visited Jan. 3, 2024) (“But the more detailed I got with the prompt requests for my replicant, the better it mimicked my tone and perspective as a journalist. The more I think about it and experiment with the custom GPT trained on my writing, the more I believe this innovation could be quite disruptive as it continues to improve.”).

<sup>116</sup> See, e.g., RICHARD SUSSKIND, *TOMORROW’S LAWYERS: AN INTRODUCTION TO YOUR FUTURE* 231–32 (Oxford Univ. Press 3d ed. 2023) (“The unspoken truth is that the current set-up generally results in clients paying for the training of law firms’ aspiring young lawyers. . . . [T]his delegation has been one mainstay in supporting the pyramidal model of profitability that has enjoyed such unchallenged success until recently.”).

<sup>117</sup> See, e.g., *id.* at 263 (“In truth, for much of the legal market the current model is not simply unsustainable. It is already broken.”).

<sup>118</sup> E.g., *id.* at 12 (“[C]lients from major companies and financial institutions are facing the prospect of an increasing workload and yet diminishing legal resources. Something surely has to give here.”); *id.* at 243 (noting, “clients are increasingly asking their legal advisers to find ways of reducing the costs of routine and repetitive work”).

<sup>119</sup> “Two possible consequences might follow for young lawyers. The first is that, with the exception of the supremely talented, young lawyers might come to get paid less in their early years of working in law firms

Second, another aspect of the large-firm career model that will be affected by GenAI is due diligence and document review work. Here also, the idea is that junior lawyers “learn on the job” through repeated immersion in tedious but deeply substantive document review. On this point, one article generously predicts:

AI does have the potential to transform the junior lawyer’s role substantially. No longer will a junior attorney be stuck sorting through documents in a deal room. Now, at the very start of their career, a lawyer can become a valuable part of innovative, rewarding client services.<sup>120</sup>

While this sounds great, it assumes that junior lawyers will have the substantive practice-area skills to do this high-level work out of the law school gate. This assumption is dubious, and it is not clear where junior lawyers will obtain those needed skills.<sup>121</sup> The traditional law school curriculum does not provide them. Perhaps legal education will change, or perhaps some sort of low-paid “residency” will emerge in which the early-career lawyer is taught the substantive skills needed to bridge the education-practice gap, as has happened in the medical profession.<sup>122</sup>

Regardless, we agree that economic forces will increasingly push the due diligence grunt work of document review and summarization to GenAI.<sup>123</sup> This transition has already been ongoing for a while, as there are plenty of companies out there who advertise their

---

than in the past. The second, and more likely result, is that law firms will recruit young and aspiring lawyers in smaller numbers.” *Id.* at 232.

<sup>120</sup> Chris O’Leary & Raees Nakuhuda, *How AI for M&A Due Diligence is Changing Every Aspect of the Deal Process*, THOMSONREUTERS, Apr. 16, 2023, <https://legal.thomsonreuters.com/en/insights/articles/how-ai-and-document-intelligence-are-changing-the-legal-tech-game#can-due-diligence-truly-be-automated> (last visited Jan. 3, 2024).

<sup>121</sup> SUSSKIND, *FUTURE OF LAWYERS*, *supra* n.116, at 230 (“If the basic legal work, upon which young lawyers used to cut their teeth, is to be undertaken by others, how does a young lawyer take the early steps towards becoming an expert? . . . [T]his is an obstacle with which very few clients have sympathy.”).

<sup>122</sup> *See, e.g., id.* at 233 (wondering whether “a reversion to some variant of the apprenticeship model” is at hand).

<sup>123</sup> *E.g., O’Leary & Nakuhuda, AI for M&A Due Diligence*, *supra* n.120 (“Incorporating AI into due diligence may present a career challenge to a junior lawyer. After all, up until recent years, much of an entry-level legal professional’s job has entailed grappling with document collection, organization, and analysis. AI will indeed automate many of these tasks, increasingly making some of this clerical work redundant.”).

AI tools for document review.<sup>124</sup> GenAI stands to accelerate this trend, but with a risk-reward tradeoff that will shift with time and circumstances, because due diligence work is very different from first draft work in light of the critical rule: *No task of material consequence can be farmed out to an unsupervised GenAI.*

With first drafts, the senior lawyer is in a strong **position to supervise** the GenAI, because the senior lawyer knows both the outcome (X document that causes Y legal posture for Z client), and the data the GenAI is drawing upon (other, similar documents that caused Y legal posture to occur). And the senior lawyer will review the draft before anyone relies on it, detecting hallucinations and mistakes, and thereby **controlling the consequences**. With both tests met, farming first-draft work out to well-trained GenAI models with access to lots of reliable sample text seems very efficient.

But due diligence and document review work are different. Here, the senior lawyer does not know the preordained outcome—that is the whole point of performing due diligence. And the possibility of hallucinated answers cannot be eliminated, no matter how well-trained the GenAI model.<sup>125</sup> Supervising the GenAI here is quite difficult.

Of course, supervising junior lawyers who perform due diligence work is also difficult. Human beings too can make mistakes. The result here becomes a risk-balancing, cost-benefit exercise around **controlling the consequences**. “Deciding who decides”<sup>126</sup> (human or machine) will become an increasingly important part of a senior lawyer’s duties. This is a legal-tactics decision made by the senior lawyer in light of the client’s risk tolerance level, which is in turn a legal-strategy choice that primarily belongs to the client. And just as a senior lawyer must know the capabilities of junior lawyers available for due diligence work, the senior lawyer must also know the strengths and weaknesses of GenAI for that same work. Here the tradeoff of what tasks can be given to probabilistic machines resurfaces, as do the challenges discussed above:

---

<sup>124</sup> E.g., <https://www.revealdata.com/> (last visited Jan. 3, 2024).

<sup>125</sup> There are emerging techniques to reduce, but not eliminate, the occurrence of hallucinations. See, e.g., Kyle Wiggers, *Are AI Models Doomed to Always Hallucinate?*, TECHCRUNCH, Sept. 4, 2023, <https://techcrunch.com/2023/09/04/are-language-models-doomed-to-always-hallucinate/> (last visited Jan. 13, 2024) (quoting an expert as saying that LLMs “do and will always hallucinate”).

<sup>126</sup> See SHOSHANNA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM: THE FIGHT FOR A HUMAN FUTURE AT THE NEW FRONTIER OF POWER* 181 (Hachette Book Group 2019) (asking, “Who decides who decides?”).

- What are the “Black Swan” problems that a GenAI may miss, where low-probability but high-consequence issues might be lurking in the documents? What “Black Swan” problems might GenAI introduce (by hallucination or otherwise)?
- What are the “Sorcerer’s Apprentice” problems that a GenAI may be glossing over, allowing individually low-consequence errors to pile up into a high aggregate-consequence issue?
- What is the materiality threshold for these risks where I (or the client) will need to accept the higher cost of human intellectual labor to address them?

These questions do not have easy answers or formulaic approaches. Like all fact-based balancing tests, the answer will be, “it depends.”<sup>127</sup> Happily, emerging research shows that, once a human user has some substantive knowledge, AI tools can help mid-level knowledge workers take big steps up in their productivity.<sup>128</sup> Hybrid GenAI / human teams are likely, where humans “spot check” the work of GenAIs to reduce risk.

Finally, we see the line between “pure” legal businesses and “pure” software businesses growing ever hazier. It was one thing when the primary capabilities of AI tools were classifying photos as “cat or dog.” But the textual capabilities of GenAI open up a new front on the long-staid legal industry’s document-based business model. Software tools created by both legal outsiders and insiders and based on GenAI underpinnings will increasingly pervade the legal services sector. As one leading commentator notes:

---

<sup>127</sup> The problem is well described by Peter Bernstein: “The story that I have to tell is marked all the way through by a persistent tension between those who assert that the best decisions are based on quantifications and numbers, determined by the patterns of the past, and those who base their decisions on more subjective degrees of belief about the uncertain future. This is a controversy that has never been resolved.” BERNSTEIN, *AGAINST THE GODS*, *supra* n.65, at 6.

<sup>128</sup> *E.g.*, Simon Johnson: *Achieving Shared Prosperity in the Age of AI*, UC IRVINE CENTER FOR DIGITAL TRANSFORMATION, Nov. 6, 2023, <https://youtu.be/deCdgDo0Jbo?t=2568> (Johnson: “If we had ten studies on that we’d be on our way. We have one study right now . . . If you provide customer service reps with an AI assistant . . . the less experienced people become better at what they are doing . . . .”); Jonathan H. Choi *et al.*, *Lawyering in the Age of Artificial Intelligence*, Minn. Legal Studies Research Paper No. 23-31, available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4626276](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4626276) (last visited Jan. 15, 2024) (conducting a controlled study and finding that GenAI assistance was most helpful to the lowest-skilled law students, while vastly improving the time efficiency of all).

The professions are unprepared. Many companies are still focused on selling the time of their people, and their growth strategies are premised on building larger armies of traditional lawyers, auditors, tax advisers, architects, and the rest. The great opportunities surely lie elsewhere—not least in becoming actively involved in developing generative AI applications for their clients.<sup>129</sup>

This prediction has already borne fruit. One international law firm has already worked with both a “lawtech” startup and a large incumbent tech company to co-develop a contract negotiation tool that it is selling to certain clients. While details are scant, it seems the tool is meant to enable in-house lawyers at those client companies to negotiate contracts by themselves more efficiently.<sup>130</sup> The firm is selling this tool as a software subscription service.<sup>131</sup>

This seems like a model for what is to come.<sup>132</sup> The economic benefits of being “valued like a tech company”<sup>133</sup> will invade the legal services space. Developing technology tools will be part of the core business of law. “Law firm technology” will one day mean much more than support-level IT departments who facilitate selling the billable hours of subject matter experts. While a host of ethics challenges will arise around who can be an investor

---

<sup>129</sup> Richard Susskind & Daniel Susskind, *Generative AI Will Upend the Professions*, FINANCIAL TIMES, June 17, 2023.

<sup>130</sup> Cristina Criddle & Suzi Ring, *Allen & Overy Rolls Out AI Contract Negotiation Tool in Challenge to Legal Industry*, FINANCIAL TIMES, Dec. 21, 2023 (“The tool, known as ContractMatrix, is being rolled out to clients in an attempt to drive new revenues, attract more business and save time for in-house lawyers.”).

<sup>131</sup> *Id.* (“A&O would not detail specific financial terms around the contract negotiation tool but said clients would pay an annual subscription fee per license . . .”).

<sup>132</sup> Given the facts reported by the Financial Times, this model seems especially clever. A software product made by outside counsel and sold to in-house counsel sidesteps a number of thorny issues. Among them: Is the GenAI tool “practicing law” by itself? Who is responsible for deciding whether or not to accept the tool’s outputs? Does the decider have the right subject matter knowledge to make this choice? Etc.

<sup>133</sup> *E.g.*, Shawn Low, *On Tech Company Valuations*, MEDIUM, Oct. 31, 2022, <https://medium.com/@chshawn.low/on-tech-company-valuations-776547c2e6e3> (last visited Jan. 3, 2024) (“Tech companies are typically valued at 10–20x revenue (sometimes way more), compared to most non-tech companies which are 1–3x.”).

/ owner of these hybrid law firm / technology firm chimeras,<sup>134</sup> the economic incentive is too great for these forces to be kept at bay.<sup>135</sup>

And just as the software companies themselves moved from selling single-purchase software licenses to “X-as-a-service” subscriptions,<sup>136</sup> so also will these hybrid law / tech firms move away from selling single-purchase billable hours and towards selling recurring-revenue software subscriptions. The benefits are too compelling to resist:

Service businesses are easier to get going without investment, but they are very difficult to keep growing forever. It’s almost impossible for a law firm or marketing agency to go from 10 employees to 100, then 1000, and keep growing.

Most software companies with 100 customers have the potential to get 1,000 customers or 10,000 or 1 million customers someday—without adding the same proportion of employees. Investors can see the possibility of you getting 10x or 100x bigger with big profits.<sup>137</sup>

To wit, even some automakers are now seeking to increase their revenues by selling subscription services for physical things (like heated seats) that were formerly single-time purchases.<sup>138</sup> A bevy of legal service software subscriptions for GenAI tools developed by / with / through your favorite law firm surely cannot be far behind.

---

<sup>134</sup> See MODEL RULE OF PROFESSIONAL CONDUCT 5.4(d) (“A lawyer shall not practice with or in the form of a professional corporation or association authorized to practice law for a profit, if . . . a nonlawyer owns any interest therein . . .”).

<sup>135</sup> See, e.g., Greg Head, *5 Reasons Why SaaS Companies Are Valued on Multiples of Revenue, Not Profit*, PRACTICALFOUNDERS, June 15, 2022, <https://practicalfounders.com/articles/5-reasons-why-saas-companies-are-valued-on-higher-multiples-of-revenue/> (last visited Jan. 3, 2024) (noting, *inter alia*, that “recurring revenue is worth a lot” and “software companies have higher margins”).

<sup>136</sup> E.g., Paayal Zaveri, *Microsoft, Salesforce, and Adobe Have Ridden to New Highs by Charging for Their Software Every Month. Here's How the 2008 Financial Crisis Helped Make Subscriptions the New Normal in the Software Market*, BUSINESS INSIDER, Dec. 19, 2019, <https://www.businessinsider.com/subscription-economy-salesforce-microsoft-adobe-explained-2019-11?op=1> (last visited Jan. 3, 2024).

<sup>137</sup> See, e.g., Head, *5 Reasons*, *supra* n.135.

<sup>138</sup> E.g., James Vincent, *BMW Starts Selling Heated Seat Subscriptions for \$18 a Month*, THE VERGE, July 12, 2022, <https://www.theverge.com/2022/7/12/23204950/bmw-subscriptions-microtransactions-heated-seats-feature> (last visited Jan. 3, 2024) (“Carmakers have always charged customers more money for high-end features, of course, but the dynamic is very different when software, rather than hardware, is the limiting factor.”).



## 6. CONCLUSION

*“Knowledge has consequences, not always intended, not always comfortable, not always welcome.”*

– Richard Rhodes <sup>139</sup>

In the wake of IBM’s “Deep Blue” defeating Gary Kasparov in a chess match in 1997, it became fashionable to predict the end of human reasoning.<sup>140</sup> The same hand-wringing emerged in 2015 when Google’s “AlphaGo” program defeated the world’s best Go players. And the emergence of GenAI has once again ignited the embers of AI hyperbole, heralding the fall of human reasoning supremacy. The end is nigh; can “Chief Justice Robots” be far away?<sup>141</sup>

We see things very differently. We view GenAI as part of a long-running trend towards human-machine synthesis.<sup>142</sup> GenAI has now developed powerfully competent “System 1” reasoning that can create convincing “shoot from the hip” level responses faster and more completely than any human can. But GenAI lacks any “System 2” comprehension capabilities, making it unable to reason about whether its response is even true—let alone in tune with the spirit of the prompt. Therefore, *it is the user who must supply the System 2 thinking that GenAI lacks.*

The future thus lies with those who can fuse the System 1 power of GenAI with their own System 2 reasoning skills to create a human-machine team.<sup>143</sup> In his work *Range*, author David Epstein calls these human-machine hybrids “centaurs.”<sup>144</sup> And for all the hype

---

<sup>139</sup> RICHARD RHODES, *THE MAKING OF THE ATOMIC BOMB* 4 (Simon & Schuster 2012 paperback ed.).

<sup>140</sup> *E.g.*, DAVID EPSTEIN, *RANGE: WHY GENERALISTS TRIUMPH IN A SPECIALIZED WORLD* 22 (Riverhead Books 2019) (noting that Kasparov’s match with Deep Blue was “billed as the final battle for supremacy between natural and artificial intelligence”).

<sup>141</sup> *Cf.* Eugene Volokh, *Chief Justice Robots*, 68 DUKE L.J. 1135 (2019), available at <https://scholarship.law.duke.edu/dlj/vol68/iss6/2> (last visited Jan. 2, 2024).

<sup>142</sup> *E.g.*, Janik Festerling & Iram Siraj, *Anthropomorphizing Technology: A Conceptual Review of Anthropomorphism Research and How It Relates to Children’s Engagements with Digital Voice Assistants*, 56 INTEGRATIVE PSYCHOLOGICAL AND BEHAVIORAL SCIENCE 709, 710 (2022) (noting “the human-technology dyad which has evolved since humans first used stones to break open coconuts or fallen trees to bridge rivers”), available at <https://doi.org/10.1007/s12124-021-09668-y> (last visited Jan. 2, 2024).

<sup>143</sup> *See, e.g.*, EPSTEIN, *RANGE*, *supra* n.140, at 23 (“Kasparov concluded that the humans on the winning team were the best at ‘coaching’ multiple computers on what to examine, and then synthesizing that information for an overall strategy.”).

<sup>144</sup> *Id.* at 23.

about the superiority of AI systems, at least in chess, centaurs have trounced even the most powerful supercomputers.<sup>145</sup> It turns out that humans and machines together are more powerful than either is alone, but it takes a shift in thinking that can be hard for expert-level human beings to swallow:

In the end, Kasparov did figure out a way to beat the computer: by outsourcing tactics, the part of human expertise that is most easily replaced, the part that he . . . spent years honing.<sup>146</sup>

Many attorneys (us included) have spent years honing the tactical skills of lawyering life: memorizing legal magic words and phrases, sharpening encyclopedic knowledge of statutes and cases, and cultivating a vast bank of “pro-X” and “pro-Y” documents and clauses. Alas, this textual smorgasbord is the very thing most easily replaced by machines.<sup>147</sup> If GenAI can create a Van Gogh style lithograph of a cat playing a banjo in Western wear parachuting from the sky, surely it can create a pro-purchaser acquisition agreement for a leased-up multifamily property styled under your favorite model form.

But *strategy* is beyond machine ken as of this writing. While GenAI models will continue to advance on their own as tech companies pour resources into them, the abilities of the human operator will continue to be the key differentiator: The better tuned the System 2 skills of the user, the more powerful the human-GenAI centaur. And because GenAI is controlled through natural language prompting, those who possess the most powerful natural language skills and a deep subject matter knowledge of the issue at hand are those best positioned to wield GenAI models. These are the core skills of “thinking like a lawyer.” They can be applied to reduce hallucinations *ex ante* and detect them *ex post*, increasing the savvy lawyer’s operating efficiency and leverage.

And so, far from being a threat to the existence of lawyers, we see GenAI as a powerful catalyst for the best System 2 abilities of lawyers. In our view, the most important thing that GenAI may generate is more lawyers—or at least more people who must learn to “think like a lawyer.”

MJT & JLE

---

<sup>145</sup> *Id.* (“A duo of amateur players with three normal computers . . . destroyed Hydra, the best chess supercomputer . . .”).

<sup>146</sup> *Id.* at 24.

<sup>147</sup> *Id.* at 22 (“Thanks to their calculation power, computers are tactically flawless compared to humans.”).