

Ask A Mentor: How Can I Promote Thoughtful Use Of AI?

By **Eran Kahana** (March 20, 2025)

Experts answer questions on career and workplace conundrums in this Law360 Pulse guest column series. Have a question you're afraid to ask your law firm chair, practice area leader or mentor? Submit it anonymously here.

In this installment, Eran Kahana at Maslon offers tips for spotting texts created with artificial intelligence and implementing AI within a firm.

My firm has started implementing some use of artificial intelligence technology into our workflow, and I'm excited about it as a potential time-saver, but I want to make sure it's only being used for the proper purposes. How do I identify AI-generated work product, and then thoughtfully address the issue?

— Partner at BigLaw firm

Similar to U.S. Supreme Court Justice Potter Stewart's simple and pragmatic "I know it when I see it" declaration about obscenity in his 1964 concurrence in *Jacobellis v. Ohio*, AI-generated content often has qualities that, while sometimes resisting formal definition, tend to become increasingly recognizable the more you are exposed to it. As I mention to my law students at the University of Minnesota Law School, "if you use AI, I'll know it."

This, by the way, does not mean my students are not allowed to use it. It just means that they need to use it properly. You will have a clearer understanding of this shortly.

But before we go any further, it is important to pause and appreciate that these "I know it when I see it" characteristics will gradually vanish as AI gets more intelligent, which it does. Every day. And when those qualities do vanish, it should be viewed as an inflection point where the question of what constitutes AI-generated is no longer relevant, simply because it is good enough, or even better than what humans can generate.

With that, we now turn to the question of how to identify AI-generated content for the time being.

Identifying AI-Generated Content

One of the most obvious characteristics is a canned, clichéd writing style that reeks of superficial engagement with the topic. It is peppered with all the popular buzzwords as if it were a shroud for legitimacy and coupled with predictable phrasing. Let's take a look at how this pans out.

In the following example, which I generated with Anthropic's Claude 3.7 Sonnet,[1] you will see a narrative containing AI-generated indicators like spots on a leopard:



Eran Kahana

In the ever-evolving landscape of legal technology, the integration of artificial intelligence presents both unprecedented opportunities and profound challenges. As courts and practitioners grapple with the implications of algorithmic decision-making, it becomes increasingly apparent that the traditional frameworks of jurisprudence must adapt to accommodate these novel paradigms. While AI promises enhanced efficiency and access to justice, stakeholders must remain vigilant against the erosion of due process principles that form the cornerstone of our legal system. The black-box nature of certain AI applications may fundamentally undermine the transparency and reasoned decision-making that lies at the heart of judicial legitimacy.

As you critically think about what you just read, you will likely find that it offered virtually nothing. Zero calories. No insight nor any meaningful analysis.

Once you get a feel for what it looks like, even without knowing precisely what that "it" is, you might also discern the stitching: the seams that connect a series of independent sentences in an obviously mechanical way. Additionally, AI narrative often lacks appreciation of nuances that apply to the practice of law and is merely a product of plugging in law-related terms.

Another giveaway is the generic quality of the narrative. It lacks spirit. There's no flair, zero creativity. This is also perceptible from the distinctive cadence.

Another detection option is to throw the content into an AI detection tool, such as GPTZero or Copyleaks. Keep in mind, however, that the accuracy of these tools is all over the map, and they tend to do poorly when analyzing output generated with newer large language models. Because these newer AI models are widely and easily accessible, the likelihood of someone using older models is low to nonexistent. In other words, there is little point in using these tools.

Finally, just ask. However, it is worthwhile noting that there is such a thing as a clandestine use of AI. Even if the firm permits the use of AI, some people might prefer to keep their use of it quiet. So, asking may not provide the input one hopes for. But it's worth a try.

Ultimately, all the AI-generated characteristics mentioned above will have a short shelf life, simply because AI is getting better. AI reasoning models, such as OpenAI's o1 and o3-mini, and Claude 3.7 Sonnet, take more time to think through their responses, but they can produce output that does not suffer from the generic, canned, clichéd and other insufferable attributes.

And so, examining content boils down to using your professional judgment. If it is good, does it matter that it is AI-generated?

Addressing AI Use Within the Firm

Now we turn to the question of thoughtfully addressing the use of AI in the firm. This has to do with developing an appropriate policy and procedures.

Tailoring Your Policies

One of the most important things to start with here is to just say "no" to using boilerplate as the final product. A truly useful policy and the attendant procedures that will help engender a diligent approach and appropriate culture regarding the use of AI cannot be

accomplished through boilerplate.

Boilerplate is a document riddled with pedantic, tedious protocol. It may provide a useful starting point, but it is no substitute for meaningful, practical guidance.

Testing Trustworthiness

The next item on the agenda is ensuring that the AI application your firm is signing up for is trustworthy. This information does not come from the marketing materials. Unsurprisingly, all we see there is how life-changing the application will be and other grandiose declarations.

Instead, the most important information is found in the terms and conditions. And that's where it is virtually guaranteed you will encounter a problem.

In the terms and conditions, you will find that the developer essentially says it has no idea if the application even works. You will see that disclosed in the disclaimer of warranties, among other places. Unless you can successfully negotiate better language, which may be possible, you will need to think about the compensating measures you will need to take.

This means you will need to find ways to build your own trust in the application. That typically means using it in a sandboxed environment — that is, not using the application on live client matters until it has proven itself.

In my law school class, for example, we take a similar approach. Students select an AI application that they are comfortable with and test its performance by asking it questions on material we covered earlier in the semester. This way they can gauge whether the application's performance is satisfactory.

As the weeks go by, students get a better feel for whether the application is trustworthy. Each week, we run this exercise; the students assign a grade to their application. It is rare for it to score an A.

Whether you follow a similar process or come up with a different one, make sure it is reflected in your procedures. This gives it a better chance of becoming a standard operating procedure as new AI applications are onboarded, rather than a one-off effort that does not reflect well on the firm's diligence.

Identifying Tasks for AI and Training Attorneys on Ethical Use

Another aspect to cover in the procedures is which tasks benefit from AI use. If it's research, then explain how that is to be conducted and provide periodic training.

Part of the training should be broader than just how to use AI in legal research. Consider expanding it to a review of cases where lawyers got sanctioned for improper use, starting with the first-ever such reported case, *Mata v. Avianca Inc.* in the U.S. District Court for the Southern District of New York in 2023, then moving on to others.

In *Avianca*, the plaintiff's attorneys used ChatGPT to identify legal precedent in support of their motion to oppose Avianca's motion to dismiss. But there was one glaring problem: The cases did not exist. The court found the plaintiff's counsel acted in bad faith and violated Rule 11 of the Federal Rules of Civil Procedure by not verifying the accuracy of the citations and relying exclusively on ChatGPT.

Another useful resource to include in the training is the American Bar Association's Formal Opinion 512, the first ABA opinion that directly addresses the use of generative AI. When you present cases like Avianca alongside Formal Opinion 512, make sure to focus on how AI should be used in research, and caution on its propensity to hallucinate, or produce nonsensical but legitimate-looking content.

During training, highlight the ethical legal obligations, including the legal obligations under Rule 11 of the Federal Rules of Civil Procedure, which requires attorneys to make a reasonable inquiry into the facts and law before signing legal documents. Then talk about the consequences of failing to meet these obligations. Between Formal Opinion 512 and these generative AI cases, you will have what you need for proper training.

Ensuring Good Data Stewardship

Another part of a diligent approach to the use of generative AI comes down to your firm's data governance practice.

Think about data from a supply chain perspective. This approach, also referred to as an AI data stewardship practice, helps ensure you have the necessary controls over the data that is used in your AI rollout. It is essentially a set of principles that provides guidance on responsibly collecting, managing and using data for training AI language models.[2]

Ensuring good data governance is an important task because the quality of the data means everything. It is the lifeblood of the application; it will only be as good as the data it is trained on.

This focus was glaringly missing in *Wadsworth v. Walmart Inc. and Jetson Electric Bikes LLC* in the U.S. District Court for the District of Wyoming. Similar to Avianca, the Wadsworth attorney used his firm's in-house AI platform to generate nine cases in support of his motion in limine, eight of which did not exist.

That attorney's negligent conduct led to a fine and sanctions last month. Though the fine was fairly small, the resulting bad press was much worse, especially since the hallucinated cases were identified in early February, well after the mid-2024 release of the ABA's Formal Opinion 512.

Assessing Quality and Effectiveness

As your firm starts using AI applications, keep a watchful eye on quality and effectiveness. It will not take long before it is apparent if it is truly useful, only marginally so or a complete waste of time.

If you approach its adoption with a diligent mindset, however, it is likely that it will prove useful in the long run. Maybe not in everything you thought it would, but that is OK too.

Eran Kahana is counsel at Maslon LLP. He is a research fellow and CodeX fellow at the Stanford Center for Legal Informatics. He is also co-author of "The Law of Artificial Intelligence and Smart Machines."

The opinions expressed are those of the author(s) and do not necessarily reflect the views

of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

[1] The prompt was: "Give me an example of a canned, generic and cliched statement in a legal analysis that shows the problem of relying on AI."

[2] An example of such a framework can be found on the Stanford Law School blog
R <https://law.stanford.edu/2023/03/09/a-data-stewardship-framework-for-generative-ai/>.