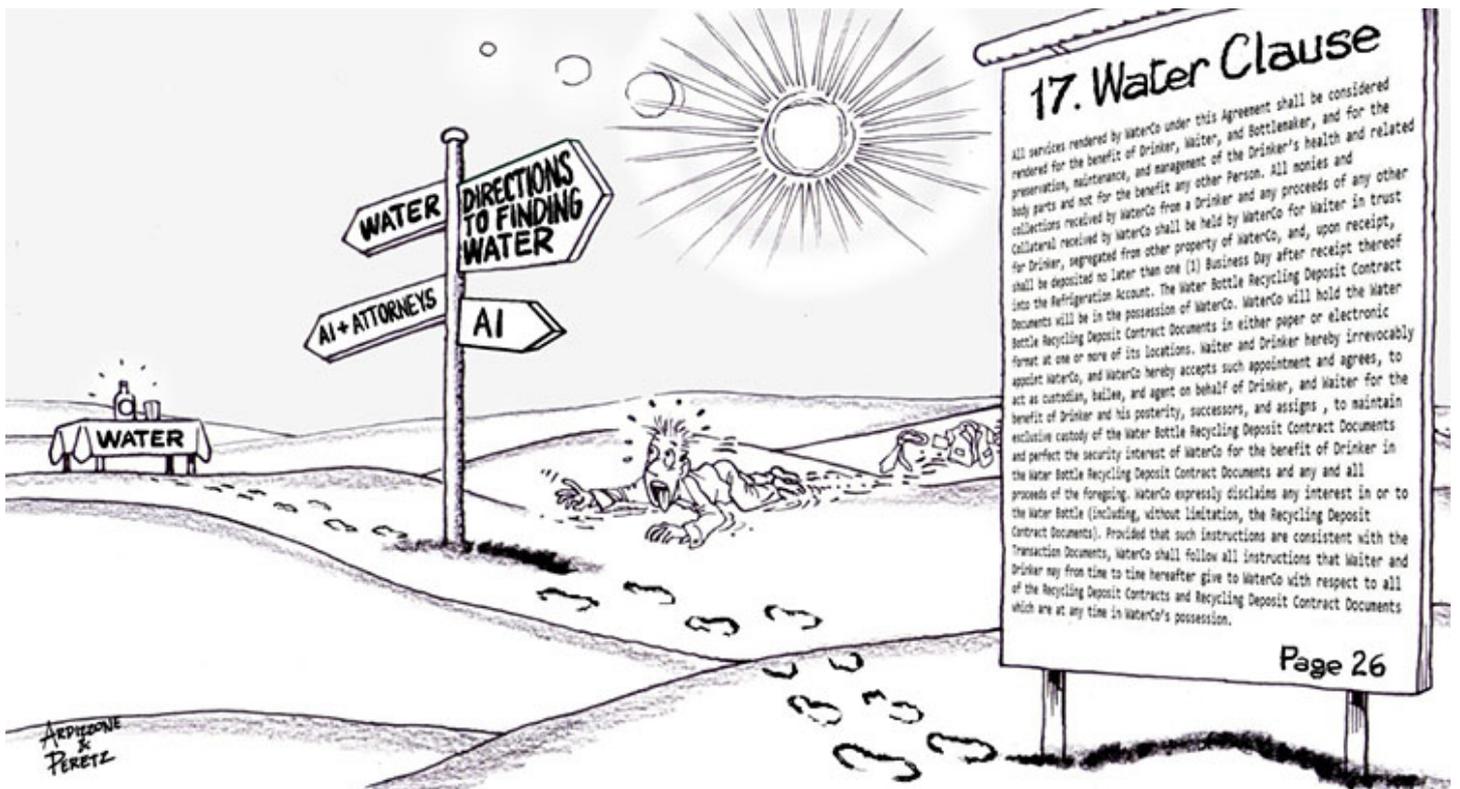
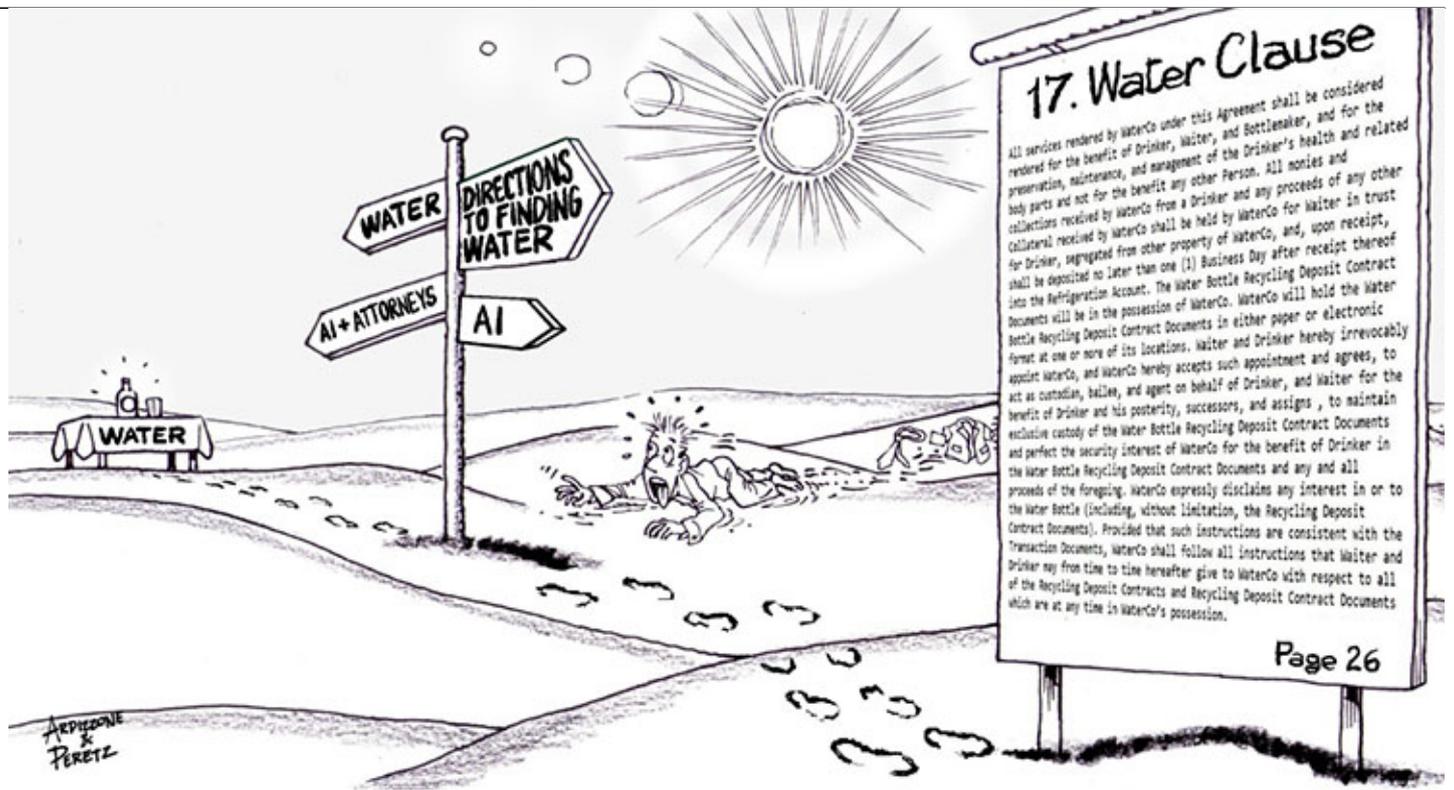




5 Questions to Ask to Evaluate AI Solutions

Technology, Privacy, and eCommerce



Art by F. P. Ardizzone. fpaulardizzone@gmail.com

As in-house counsel, we are bombarded with warnings that the machines are coming to take our jobs and new software includes artificial intelligence (AI) that will handle key legal tasks for us. Should we be scared? Or perhaps be rejoicing?

Whether you're having a nightmare about impending job loss due to the rise of AI or a fantasy about

the prospect of machines reading those lengthy paragraphs of capital letters to determine the real limit of liability in an agreement, it's time to wake up.

To avoid getting lulled into a haze again about the importance of attorneys when someone promises a quick AI solution, ask these key questions:

1. Does it classify clauses or does it provide hard, actionable answers?
2. Was the system trained by large document review teams who were under time pressure to review large quantities of documents quickly?
3. What happens when I apply the AI to agreements more complex than the sales template we use repeatedly?
4. What does the system do when an agreement has interdependencies between terms?
5. What attorney and paralegal resources are required to determine accuracy and completeness of AI results?

To answer these questions honestly, we need to learn about the limits of AI in the legal context.

Warmer, warmer, colder, warmer

A popular party game for kids and adults alike is to blindfold someone or put them in an unfamiliar setting with the mission of finding a hidden object, such as a thimble in old days. The person who hid the object provides general hints about proximity to the target phrased in terms of temperature.

When the treasure seeker gets closer to the target, she is told that she is “warm” or perhaps even “hot” if the target is very close. Similarly, when the seeker moves away from the target, he is told that he is getting “cold” and “colder.”

Today's application of AI to legal documents resembles the aforementioned Hunt the Thimble game. The role played by today's AI tools is to give directional hints regarding the sought after information and some qualitative measure your proximity to the target, essentially telling you whether you are getting “warmer” or “colder.”

Of course, they make this sound more businesslike and professional by calling it “Clause Classification” rather than “Hunt Down the Payment Provisions” or “Hide and Seek the Liability Limits.” What the AI does not do is to find and uncover the treasure for you — that task is still up to you.

Today's AI tools for contract analysis typically provide suggestions about where a particular piece of information might be found within a contract, acting like a voice whispering in your ear, “Hey, look over here and I think you'll find the Payment Terms clause.” It remains the responsibility of the attorney to read the Payment Terms section and ascertain when invoices need to be sent and how long to wait before payment is deemed overdue.

It's great to have a tour guide who can bring us to a relevant section of an agreement; however, the most difficult part of the task still remains: reading the particular section of the agreement and identifying the specific terms and actions that flow from the section, such as “Send the invoice after service delivery” and “Expect payment within 15 business days.”

Harder still is when the answer is spread across multiple clauses, like when the invoice date depends on XYZ (e.g., the second signature date) in Attachment A. AI won't pull those pieces together for

you.

The next time you evaluate an AI tool for interpreting agreements, delve into the specifics of what it helps you do and whether subsequent attorney work is required in order to derive actionable, standardized data (not contract prose) that enables the desired analysis across multiple agreements.

There's no doubt that you can sometimes "Find the (Clause)" faster with advice about the general vicinity where you should look (page 23, paragraph 9), but don't expect the attorney's role to diminish greatly because it remains your responsibility to read the relevant clause and determine how it translates into data and corresponding actions that your organization needs to take.

When planning your budget and setting colleagues' expectations for the legal department's throughput, you need to remind colleagues that technologies like AI might speed your work somewhat by finding some relevant clauses. However, you still need substantial attorney hours to take the AI suggestions and accurately translate those into specific, actionable data.

Rushed "teachers" are not the best educators

Why can't the AI go beyond telling you which section of the agreement to read next? One reason is that the traditional mechanisms for training AI are difficult to adapt to the legal context due to special characteristics of attorneys and the practice of law. Understanding legal documents requires quality even more than quantity.

You have likely been asked online to verify that you are a human by solving a CAPTCHA puzzle, such as identifying a streetlight or crosswalk in a photo. Most humans have the skills to perform such a task accurately because recognition of traffic signals is widespread.

By contrast, asking the general public to identify whether a force majeure clause includes acts of government is not likely a task well-suited for the general public. Instead, we need highly trained attorneys to answer questions about information contained in legal documents like business agreements.

Unfortunately for those who want to train AI on legal documents, highly trained attorneys are quite expensive and often committed to higher priority initiatives than leveraging their experience to train AI models. Due to their high cost, people who hire attorneys to review documents are typically most concerned with throughput and reducing this cost as much as possible.

To manage costs, attorneys are judged based on their reading and annotation speed. The attorneys most likely to stay employed on document review projects are those who are the fastest.

A focus on speed may be essential for keeping document review projects within budget, but it is not the right approach for training AI. Imagine if your teachers in school were compensated based on how quickly they could grade your papers or get through a lecture: Corners would get cut.

Suddenly the grades on one student's torts exam might be higher due to neater handwriting, rather than presentation of quality thought, because the professor rewards quick readability. And papers graded first might receive a substantially different grade than papers graded last, despite having similar content because the professor loses "speed points" by comparing his work across the entire class to normalize it.

Training AI is a lot like training students. If a document is harder to read or more complex, more time needs to be spent breaking it down and teaching the key elements to the AI software. And the human expert trainers need to constantly revisit previously reviewed documents to normalize how they are parsed to match later documents.

Rewarding extra time spent breaking down complexity and constantly revisiting prior documents is anathema to how large-scale document reviews by attorneys are conducted today due to the high cost of the attorneys and the need for speed, perhaps to hit a merger review target, court date, or other time-sensitive event.

What this means is that you need to examine how each AI system has been taught. When you hear that a system has been trained through deployment on giant document review projects, this is likely a sign that it had “rushed teachers.”

The care and nuance applied by those teachers and the results reached can vary wildly. The AI is going to bring the same wildly varied results to its analysis of your documents, which means the system is relying on your intensive review and adjustment of the results to achieve the quality standards that you and your clients are seeking.

Imagination is clean and straightforward, but reality is messy

When people imagine training a computer to read a legal document, they often conjure the image of the perfect legal document: an agreement in which each piece of information is in the right place in a clear, self-contained document. Unfortunately, that is rarely how contracts are written.

Consider a lease agreement that states: “This lease will expire on December 12, 2017.” It’s the example of a perfectly clear agreement. Surely it should not be too hard to train a computer to look for the word “expiration” in proximity to a calendar date and have a high likelihood of finding the expiration section of the agreement.

Moving from our dream contract to the real world frequently compounds the challenge. For example, what happens when the agreement states elsewhere that “If not spent, the tenant improvement allowance provided in the lease will expire March 13, 2016”? The contract now has two sentences containing a variation of the word “expire” proximate to a date and the word “lease.”

Eventually, a deep learning model might be able to distinguish these two scenarios. However, more attorney hours are still required to train such a model. Because legal documents in the real world are rarely as tidy as the documents in our imagination, it’s common to dramatically underestimate the amount of expert attorney training required to teach machines how to distinguish between correct answers and false alarms.

Business relationships have entropy too

Like many physical systems, business relationships are also prone to increasing levels of complexity and disorder. In the imaginary world, all the terms of a business relationship might be embodied in a simple two-page agreement listing the obligations of both parties.

Reality is rarely so well organized and succinct. Often the terms of contractual relationships are spread across multiple documents, such as invoices, purchase orders, statements of work, and terms

of service posted elsewhere. Each of these adjacent documents has legal consequences for the business relationship, but they are composed in a structure and form language substantially different than the prototypical contract.

On top of the wide range of documents to cull through, some key business terms might also have a dependency that cannot be determined within four corners of any of the documents. For example, many leases commence on the earlier of a specific date (e.g., March 1, 2015) or the date that the tenant starts to utilize the space.

AI cannot figure out the lease commencement date in this example because the contract does not tell us whether the tenant move-in date preceded March 1, 2015. We still need human attorneys to recognize that a determination of the answer depends on missing information and to seek that information from external sources.

It's often a scavenger hunt, not a game of tag

Another reason that today's AI tools still require substantial human attorney input is that legal documents often encode actionable answers into multiple puzzle pieces across the document. AI may be able to parse self-contained sections of a document, but often stumbles when the ultimate answer requires finding and synthesizing many sub-answers along the way.

Consider parsing the confidentiality section of a typical agreement to determine when confidential information needs to be returned or destroyed:

Potentially suitable for AI

Section 7. Confidential Information must be returned before June 30, 2027.

Extremely difficult for AI

Section 1. This Agreement is effective upon the last signature of a party to the Agreement.

Section 3. The Expiration Date is the third anniversary of the first day of the month following the Effective Date.

Section 7. The Confidential Information must be returned within three months of the Expiration Date.

Signature blocks

Signed by Borrower Jane Doe on February 26, 2019.

Signed by Lender First Missouri Bank on February 28, 2019.

Signed by Escrow Agent as having been received on March 2, 2019.

For either agreement described above, a typical AI tool would tell the reader to review "Section 7" to determine when the confidential information needs to be returned. However, the more detailed agreement (on the right side) requires the review of many other sections of the agreement in order to convert lengthy, scattered contract prose into actionable data: That the confidential information must be returned by June 1, 2022.

Because agreements are regularly amended and expanded, many of them contain interdependencies across multiple clauses and terms. The result is that AI tools can sometimes provide a starting point for research, however, extensive attorney study of the document, as well as

related documents (e.g., the 12 Statements of Work that accompany a Master Services Agreement), is often required in order to derive actionable data.

[Neil Peretz](#)



General Counsel

Sawa Credit Inc.

Neil Peretz has served as general counsel of multiple companies, particularly in the financial services and technology industries, as well as a corporate CEO, CFO, and COO.

Outside of the corporate sphere, he co-founded the Office of Enforcement of the Consumer Financial Protection Bureau and practiced law with the US Department of Justice and the Securities and Exchange Commission. Peretz holds a JD from the University of California, Los Angeles (UCLA) School of Law, an LLM (master of laws) from Katholieke

Universiteit Leuven (where he was a Fulbright Scholar), bachelor's and master's degrees from Tufts University, and has been ABD at the George Mason University School of Public Policy.

He previously co-founded legal technology company Contract Wrangler, which applied artificial intelligence to read legal agreements. Follow him online at [linkedin.com/in/neilperetz](https://www.linkedin.com/in/neilperetz).