



# A first-hand experience with AI legal software

## An appellate lawyer's review of Casetext's CoCounsel

By **JEFFREY I. EHRLICH**

A few months ago, I learned that the legal-research vendor, Casetext, had developed a generative AI legal-research tool called "CoCounsel," built on the ChatGPT platform. I have been interested in all aspects of AI for years. Soon after OpenAI released ChatGPT as a free app in the Apple Store, and Microsoft started to incorporate it into its Bing search engine, I learned about CoCounsel. I was curious enough to sign up and try it.

After I signed up, the parent company of Westlaw, Thomson Reuters, purchased Casetext for \$650 million in cash. The sale has since closed. Westlaw will debut its new generative-AI product, called "Ask Westlaw," on November 15, 2023. That product will have legal-research capability similar to CoCounsel, but with the benefit of having been trained on all of Westlaw's content. It will also offer some, but not all of CoCounsel's other capabilities. Those will be gradually imported into Westlaw in the coming months.

I'll start with some caveats. First, I am an appellate lawyer, and the bulk of my practice involves doing research, summarizing appellate records, and writing appellate briefs. As a result, I did not have a need for some of the "skills" that CoCounsel is marketed as offering, such as creating outlines for depositions. There may be things that CoCounsel can do for your practice that I had no need for, or that never occurred to me to try. In short, what worked – or didn't work – for me, may not predict what will work for you.

Also, CoCounsel is evolving. During the time I subscribed, Casetext rolled out various enhancements. For example, when I first subscribed there was no way to limit search results to published cases.

But within a few weeks that option was added. So, it is possible that some aspects of the program have been upgraded or enhanced since I tried it.

### **What did it cost? \$500/month**

I am a sole practitioner. When I signed up for CoCounsel a few months ago, there were two pricing options available. I could either buy an unlimited subscription for \$500 per month, which renewed automatically but which could be canceled, or I could pay \$50 per query or task. I don't know if those options have changed, or how much it costs a multi-lawyer firm. I opted for the \$500 per month option and used it for three months.

### **What can it do?**

According to Casetext's marketing information, CoCounsel can review documents, allowing you to find information within them; it can help you prepare for a deposition by generating an outline of questions for you; it can search a "database" of documents you upload to it; it can summarize transcripts or other documents; and it can draft legal-research memos based on your natural language prompts.

### **What did I use it for?**

I had CoCounsel summarize a couple of transcripts. I tried to use it to help me find documents in a "database" I created, which consisted of a law firm's file for a matter that I am litigating. I used the deposition-outline feature, just to see how it worked. But mostly, I used the legal-research memo feature.

### **How well did it work, and was it "worth it"?**

CoCounsel is impressive. I can see use cases for some law firms that would

justify the investment. But after using it for 90 days, I concluded that it was not worth \$500 per month, on top of my Westlaw subscription. But when Westlaw folds it into its product, I'm likely to purchase access to it, provided that the pricing is not too prohibitive.

### ***Depo prep***

I was particularly impressed with its deposition-preparation tool. It asks you to describe the deponent and the critical issues in the case, and it generates an outline of potential questions and areas of inquiry. I thought that it seemed to cover most of the important issues. If your practice involves cases that tend to present the same issues from case to case, such as personal-injury cases arising from auto accidents, it could be a time-saver – particularly if your firm is relying on inexperienced associates to take the depositions. But because I seldom have need to take depositions, this feature was not all that useful to me.

### ***Summarizing transcripts***

I was impressed with this feature too. I uploaded a normal-sized volume of trial testimony, and in about eight minutes it produced a credible summary of what happened. Was it good enough for me to rely on to prepare an appellate brief without looking at the actual transcript? No. Not close. But if what you need to know is, in general terms, what questions were asked and how were they answered in a deposition or a day of trial testimony, this tool does that reasonably well.

Before I started my own firm, I used to handle legal writing for a larger plaintiff's firm. I would often be unaware of a case until a summary-judgment motion landed on my desk, together with a file that included several large depositions.



But no one had prepared a summary of those depositions. As a result, I would have no idea which witnesses were important and which had little to say on the issues raised in the motion. As a result, I would often have to spend several frustrating hours reading deposition transcripts, just to know that I did not have to review them.

I think that CoCounsel's summaries would be a great tool in this situation. It would have allowed me to upload the transcripts to be summarized. Then, based on those summaries, I could have quickly discerned how important that witness was to the motion. That would have saved me hours of needless deposition summarizing.

So, if your firm has lawyers taking lots of depositions, this feature would allow your staff to create summaries for all depositions very quickly, and any new lawyer picking up the file could quickly learn the gist of what the deponents said. I think this could be very valuable.

#### **Large databases – some limitations**

I was hoping to use this feature like an e-discovery platform, to help me find critical documents in the file of a law firm being sued for legal malpractice. The file had been imaged into 50,000 different PDF files. I had been using Adobe's search tool to perform keyword searches to find relevant documents. It was cumbersome and I had concerns that the search results were not all that reliable.

Uploading 50,000 PDFs into Casetext was not easy. They wanted me to upload them about a few hundred files at a time, which would have been a lot of uploads – far too many. So, I uploaded them all at once, which failed. I then divided the files into multiple batches, and that didn't work either. Ultimately, I uploaded the file into Dropbox and Casetext migrated it for me within a couple of days. (They were very helpful.)

I thought I was in business once the files had been uploaded into Casetext. Collectively, those files form a "database" whose contents I hoped to investigate

by harnessing CoCounsel's AI chops. I hoped that the ability of generative AI to deal with context would give me more accurate searches than by using simple key words. When I ran the first search on a critical term in the case, it returned 50 results. When I ran a different search, it also returned 50 results. The same thing happened when I ran a third search.

At that point it hit me: For some reason CoCounsel was limiting its response to 50 results, regardless of how many "hits" the search turned up.

When I raised this with Casetext support and asked how to raise the limit, they investigated and I was ultimately told that this was how the system was designed and that the number of results could not be increased. I was surprised, because this seemed to make this "skill" essentially worthless. Regardless of the number of actual "hits" in the database for a given search, CoCounsel was only going to return 50 results.

I ended up having to buy access to a dedicated e-discovery platform called Goldfynch. I have been most impressed with Goldfynch. It is powerful, relatively easy to use, and priced at a fraction of what competing e-discovery platforms charge for the same service. But that's a different story.

#### **Legal research memos – intriguing but mixed results**

This was the core "skill" that intrigued me, and I enjoyed using it. But I have to admit, the results were mixed. Here is how it works: You select that "skill" and give CoCounsel a prompt – basically like formulating a fairly detailed Google search. You tell it what jurisdiction or jurisdictions to consider and tell it whether or not to include unpublished cases in its results. Once you give it the prompt, it will reformulate it and ask you whether its formulation is accurate. It almost always was. In most cases, the reformulated version was more precise and clearer than my prompt. Once you are OK with how CoCounsel has framed your question, you tell it to prepare a research memo.

That process takes a few minutes. My tech-savvy friends tell me that it basically uses two techniques to draft the memo. First, it runs what amounts to a natural-language keyword search in its database and comes up with a set of results that should contain the information for the memo. It then uses its generative-AI platform to use those results to write a memo.

Sometimes, the results are great. A couple of times I used CoCounsel, it found exactly what I wanted on the first try. But that was not the norm. Usually, it would generate a memo that listed results that were relevant to my question in a broad sense, but which did not actually focus on the specific question.

For example, if my question was about whether there was personal jurisdiction given a factual scenario, the memo would list a few cases that seemed to match my facts reasonably well, but the bulk of the citations would be to cases that simply discussed personal jurisdiction. The memo would acknowledge that those cases did not actually answer my question but stated that they were broadly relevant to the issue.

How accurate were the memos? That varied. Once, it generated a memo that gave me *exactly* the answer I had been looking for, in a published California appellate opinion. I was elated. But also puzzled, because I knew from experience that the California Supreme Court had granted review in that case and had reversed the Court of Appeal's decision. This was not an isolated example. At least when I used it, it was clear that CoCounsel either did not "cite" or "Shepardize" the cases it cited, or that it tried, but did so unreliably.

At least as of a month ago, it is clear that CoCounsel cannot come close to drafting a reliable legal-research memo that does not have to be carefully evaluated for accuracy. Or, perhaps it is more accurate to say that CoCounsel will generate both reliable and less-reliable memos, and you will not be able to know which type you have received without conducting further research to verify CoCounsel's memo.



Ultimately, I found that using CoCounsel seemed to be a potential time-saver as the first step in doing legal research, with the second step being to use the results it gave me to formulate Westlaw searches. But after a couple months, I decided that even this approach did not justify paying \$500 per month extra on top of my Westlaw subscription, so I canceled. But I am looking forward to regaining access to CoCounsel's capabilities through Westlaw soon. Hopefully, Thomson Reuters will price it reasonably.

### Final thoughts

Generative AI is already remarkable. But the versions currently available are relatively new. The software is only going to get better – more capable, more accurate, and faster. And those improvements will likely happen quickly. Trying to predict what generative AI in the legal context will be like in two or three years is like looking at a mid-1980s IBM PC and trying to extrapolate to today's powerful hand-held phones and high-speed internet connectivity.

One thing that is clear now, however, is that one of the most important skills for lawyers to develop is how to formulate AI prompts. The quality of generative AI's outputs is substantially dependent on those prompts. And formulating useful prompts is not necessarily an intuitive skill. It takes practice and familiarity with the particular system you are using.

One lawyer I respect has found that the quality of CoCounsel's memos improves substantially if, after you run the first query, you follow up with asking it what the "best cases" are for your position, and then asking what the "worst cases" are for your position.

And a patent lawyer I know at a big firm on the East Coast did a test, pitting CoCounsel against summer clerks on the same legal issue. The results were not surprising: The summer clerks' memos were substantially superior to the memo generated by CoCounsel. But it took the summer clerks three or four days to write their memos, while CoCounsel generated its memo in less than 10 minutes. With time, the gap between the accuracy of the summer clerks' product and CoCounsel's product is likely to close, perhaps substantially. And the time savings will always be there.

Legal research is, in some ways, like creating a painting. It's not always clear when to stop. As lawyers, we tend to stop doing research once we are satisfied that we have found the "right" answer. But in law, there are always exceptions, or cases that can be distinguished or that use a different analytical framework. We don't devote the same amount of time and effort to every research project. Sometimes, the right answer is just easy to find. Other times, there is no obvious "right" answer, and we need to reason by analogy, casting a wide net to find cases in other areas that can be helpful.

For some types of research, CoCounsel will already reliably get you the right answer quickly. But then, so too would just running a search in Google or in Westlaw. For more complex questions, CoCounsel is not going to give you a clear, reliable answer.

Going forward, lawyers and law firms will have interesting issues to balance in deciding how accurate is "accurate enough" for a particular project, given the substantial time savings that generative AI already offers.

But AI can already do many things that lawyers do, or need done, and it can do many of those things pretty well. My only prediction is that lawyers of all stripes will have to become aware of what AI can do, and will have to learn the best ways to employ it. It probably makes good sense to get started now.

*Jeffrey I. Ehrlich is the editor-in-chief of Advocate magazine and an emeritus member of the CAALA Board of Governors. He is the principal of the Ehrlich Law Firm, APC and is certified as an appellate specialist by the California Bar. He is also the chair of the California Academy of Appellate Lawyers' Task Force on Generative AI and the Law.*



Ehrlich

