

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]



Full Episode Transcript

With Your Host

John E. Grant

[The Agile Attorney](#) with John E. Grant

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

When it comes to their clients using AI on their own legal matters, most lawyers' default will be to give a straightforward piece of advice: just don't do it. Now, I understand the instinct, but I think it's the wrong call, not because it's bad advice, but because I don't think it's a realistic expectation anymore. A good chunk of your clients are going to use AI tools whether you give them permission or not, in no small part because the tools make it easy for clients to scratch an itch that lawyers aren't always great at scratching anyway.

So today, I'm making the case that your job isn't necessarily to stop your clients from using AI. It's to help them be educated users of a toolset that is both powerful and potentially dangerous. I've got six concrete topics you should cover with your clients, plus a downloadable one-pager to go with this episode.

You're listening to *The Agile Attorney Podcast*, powered by GreenLine. I'm John Grant, and it is my mission to help legal professionals of all kinds build practices that are profitable, sustainable, and scalable for themselves and the communities they serve. Ready to become a more Agile Attorney? Let's go.

Hey everyone, welcome back. So, last week, I re-ran my voice of the client interview with Jason, and I think there's some really compelling stuff in there, so I encourage you to listen to it, but no worries if you didn't. Here's the punchline. Jason was a modest means client who sold his motorcycle to afford the \$5,000 retainer for his divorce lawyer.

And he wound up in a tough spot legally when the retainer ran out. The big tension for Jason was that he needed the lawyer for more than just the practical details of his case. He had questions and concerns that he'd hoped to get answers to, but the lawyer, correctly, told him that she'd need to charge him to answer each and every question he sent over, which put him in this tension between living with the uncertainty and running out his retainer even faster than he already was.

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

And while Jason's story definitely has an access to justice component, I think the frustration and emotion he describes are not uncommon experiences for clients across the income spectrum. Now, long-time listeners know that I have a framework for what clients truly need from a legal service.

I first laid it out way back in episode six, and it goes like this. There are two practical needs that people want to satisfy when they reach out for legal help. They need help mitigating risk, and they need help navigating complexity, full stop. But then there are also these three social emotional needs that people have, and I refer to them as wisdom, advice, and consortium. And that last one, consortium, is the sense that somebody is in this with you, that you're not carrying the thing alone.

And it turns out Jason's lawyer was doing her best to meet the practical needs, but again, correctly, knowing her client's financial situation, she was reluctant to provide for the social emotional ones. Now, I recorded that episode back in 2021, and the actual legal case happened several years before that. Back then, Jason's alternatives were a law librarian and a whole lot of free attorney consults.

But today, someone in Jason's shoes just opens a chat window with any of a dozen different AI tools, and that chat window will absorb every question he has at any hour of the night without judgment, without impatience, and without really making him feel like he's running up a bill.

So the question I want to explore today is not whether your clients are using AI in this way, they almost certainly are. The question that I want to explore is how you can incorporate that knowledge, that information into your law practice. What are you going to say to your clients about it? Because they need to hear it from you. And right now, I think a lot of lawyers are saying nothing or maybe saying some of the wrong things.

Now, before I get to the interesting part, at least from my perspective, let me deal with the legalistic stuff, right? The rules of evidence, the sort of

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

defensive things. And I recognize there are very real legal concerns when clients use AI on their own case. There's confidentiality, there's attorney client privilege, there's work product protection, there's discoverability of chat logs. And courts are actively wrestling with these questions right now. I think there have been at least three different federal decisions in 2026 alone, and they don't all point in the same direction. The state level stuff is even more diversified.

This is not a legal ethics episode, so I'm going to acknowledge those challenges, and we'll talk a little bit about them, but I'm not going to go deep on that. Also, others are covering that ground pretty well. So I'll put a link to a good deeper dive in the show notes if you want to look for it.

But here's what I'll say about the standard sort of knee-jerk advice I'm seeing out there, which is basically, tell your clients not to use AI. And that might feel good from a cover your rear end perspective, but it isn't realistic, and honestly, it may not be that helpful from the client's point of view. Your client is going to use these tools, and the only question is whether they use them informed or uninformed.

I also want to do one other little diversion before I roll into the meat of this thing. I kind of want to get curious for a minute. Welcome back, Ted Lasso, because I think that AI is meeting some real client needs, and we can learn from what people are doing with it. Maybe not to outcompete it at its own game, but at least to better anticipate and meet our client needs in ways that will make them want to use us lawyers instead of avoiding us, right? Figure out how to use the hybrid solutions.

So what are people getting from AI? First and foremost, and frankly, hardest to compete with, it's always on, right? Nights, weekends, two in the morning when anxiety spikes, there's no waiting for a call back, there's no rushed conversation. When the client's racing brain shows up, and you know the racing brain will show up, the machine is there.

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

Number two, it's pretty cheap to talk to, at least in terms of dollars. And that means your client can ask as many questions as they want. Now, whether the AI's answers are any good varies wildly depending on the tool and the context and the client scenario, but the client can access their AI effectively for free with no meter running and also no judgment from the tool. And compare that to the experience of asking questions of a lawyer for the third or the fourth time. It can start to feel like you're questioning their judgment if you're the client, and also, by the way, that just cost you another point two on your bill.

Number three, and this is kind of related, the tool lets them war game out all of these various scenarios that their racing brain is going to come up with. What if I do X? What if my opponent does Y? What if the goldfish dies? The client can run all of these hypotheticals as many times as they want, and all of those what if questions that would probably be prohibitively expensive to work through with you or maybe just feel embarrassing, they're going to do with their computer system, their AI tool.

And of course, these are all things that are totally natural and very human for your client to want to do. And AI makes indulging that instinct so seductively easy. And I do think that we as lawyers should try to find workable ways to better meet these types of client needs. Better road mapping, which I talk about in episode 108, more regular check-ins with the client, maybe some FAQs or video lessons covering common topics for people in your client's situation. But even with that, the temptation for clients to use AI to get that quick fix is going to be strong, almost impossibly hard to ignore.

Of course, what the client wants to do and what's good for them are two different things. And that gap between what they want and what actually helps them solve their legal problem is, of course, exactly where they need a lawyer and not a chat window, which brings me to the meat of today's episode. Because you're probably already starting to see the problem.

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

While your client is getting all of that availability and patience and maybe judgment-free conversation, the AI is also doing some things to them that they almost certainly don't understand and probably don't realize is happening. And this is the type of conversation you need to be having with them early in your engagement and maybe later in your engagement too. Because these AI tools are not neutral.

And I covered some of the pitfalls of AI used by lawyers back in episode 122, which is part of this same occasional AI ROI series. But today, I want to frame those same pitfalls around what's happening with your clients. Because I think there are six key things that lawyers should be talking about with their clients with respect to the client's use of AI. And this probably needs to happen both up front and during the course of the matter.

And the first one is that set of legal risks that I kind of glossed over earlier, but those are good ones to lead with because they're the ones that could genuinely take down somebody's case. But of course, this is also the thing that lawyers tend to naturally care about way more than their clients do. So the challenge is framing it in terms your client can understand and relate to.

Here's one way to think about it. Your client has probably heard stories about people losing their case because of something they posted on social media or said in a text message. So you probably should encourage people to think about AI in exactly the same way. The things they say in their AI chat may very well wind up in front of the other side.

Of course, the client is not going to want to avoid AI anymore than they're going to want to avoid social media. And they may not be able to contain themselves, but you've got to at least lay the reality out on the table. So hopefully, they can resist the temptation, but at least they won't be ignorant of the consequences if they don't.

The second thing is to inform them about the hallucination problem. And keep in mind, this may genuinely be news to your clients, even though it's

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

probably old news to us attorneys. These invented cases, misstated rules, confident answers about laws that either don't exist or maybe aren't right for the jurisdiction you're in. AI is this grab bag of useful information. It might be good information that doesn't apply to your specific situation, and it could be stuff that is just plain wrong or invented from whole cloth.

But clients are not nearly as steeped in these problems as most lawyers are. And in a way, legal is one of the only places where people have been having these public and documented consequences from poor AI hygiene. So we need to educate our clients.

Here are some things I think we need to get across. Help your client see that AI acts with all of the misplaced confidence of a typical 14-year-old arguing with their parent, which is to say AI sounds every bit as confident when it's wrong as when it's right. There's no easy tell. So any answer it gives you or your client about the case is at best a lead to be checked, not a fact to be acted on. And wouldn't your client rather have their lawyer be the one doing the checking than have them be the one cleaning up the mess of bad AI information or advice?

The third thing is to help the client understand that AI has this tendency to blow sunshine up their rear end. This is what I've discussed before as the sycophancy problem. These tools tend to tell people what they want to hear. So if your client is frustrated with how their case is going, the AI will validate that frustration. If the client's skeptical of your strategy, the AI will find reasons the skepticism is warranted, even when or especially when the frustrating advice you gave them is the right advice.

So you've got to try to convince your client that these tools are designed to agree with you, but that having a virtual yes man isn't going to do them any favors, even if it feels nice to be validated. If your AI is telling you everything you want to hear, that's not a second opinion, that's sort of a mirror, right? A truly useful advisor will tell you hard things, the things you don't necessarily want to hear, and that's one of the best things you're going to get from a good lawyer.

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

The fourth thing is the persistent engagement problem, right? This notion that the AI, by default, doesn't want people to stop working on a thing. I talked about this back in episode 122, but these tools are designed to keep people using the software because engagement for their users is what's good for the businesses behind the AI regardless if it's what's good for the people. You know how this goes, right? Every answer the AI gives opens two or three new questions. Every potential rabbit hole leads to yet another trail to follow.

What the tools never say is, "Hey, we're done here. This issue's settled. So stop worrying about it. Go watch some Netflix or something," because settling the question ends the session, and ending the session is bad for business. And so the information you need to get across to the client is the idea that a big part of what you're paying me, your lawyer for, is telling you what doesn't matter. When I say something's a dead end, that's me saving you time and money and worry. The AI won't close doors for you, I will, and that is a feature, not a flaw.

Number five, we need to warn people that AI can be something of a mean girl, meaning it has a tendency to stir up drama between clients and their lawyers. There's a few pieces of this. How your client prompts the AI shapes what the AI tells them. So if the client types, "My lawyer says X, tell me why she's wrong," the AI won't question the premise, it will produce a confident, well-organized case against you, the lawyer.

But if my client types, "My client says X, help me understand her reasoning," the AI will produce something totally different. Same tool, same facts, different answers, and most clients have no idea when they're tilting the AI in one direction or another. They think they're asking a neutral machine a neutral question, but the reality is they're steering and the machine is following.

Also, the sycophancy problem comes back around here. The AI tends to be loyal to its user. So even if you, the lawyer, tried to get your client to prompt it in a certain way that benefits you, the AI is still going to tend to come

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

around to supporting whoever is manning the keyboard, not necessarily your interest or even the client's interest. And I'll admit, this is a tough one to teach around because it happens so much in the background, it's so natural. But you can try showing your client those two prompts side by side. This is one of those cases where showing is probably better than telling.

The idea is to plant a useful seed. The next time the AI hands your client some argument against you or the strategy that you've developed, you're hoping the client will wonder whether the way they phrase the question is what produced that particular result.

And for you, the lawyer, it can be a helpful reminder too. I get it's frustrating when a client shows up with an AI generated argument for why your thinking is wrong, but understand what probably happened. They weren't researching, they were worried, and they asked the machine a leading question, and then the machine did what these machines do, attempt to validate the person. That's not necessarily a disloyal client, it is an anxious one.

The sixth and last thing I'll hit on for today, and I think this is a good one to end on for your client as well, is to help your client remember that AI has no stake in the outcome. For one, it doesn't really know how the real world works. It only knows what people have written about it.

And on top of that, it carries no duty to the client, right? No accountability, no malpractice coverage. It also has these huge knowledge gaps, right? It only knows what got typed into it. It can't see the procedural posture, it doesn't know your judge, it doesn't know your relationship with opposing counsel. Even when its answer is technically correct, it's a correct answer to this abstraction of their actual situation, and it's hard to get that exactly right.

And I think you can boil this down for the client in kind of one sentence. And it's this. If your AI gives you bad advice or a wrong answer, you're still

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

the one who has to live with it. The AI literally doesn't care because it isn't actually capable of care. And everything about these tools flows from that.

You need to help your client see that the AI can afford to be confident because it never pays for being wrong. But of course, your client does suffer consequences for following bad advice or repeating bad information. And you, the lawyer do too. That's precisely why your legal advice comes slower with more caveats, because unlike the AI, your advice is attached to a real human who is accountable for it. And you need to help your client really wrap their heads around that.

All right, so how do you communicate these things to the client? I think there are two sets of tools. One is something to build and one is a set of conversations to have. In terms of the tools to build, where should you put all this information?

Well, hopefully you've taken my advice from episode 108 and elsewhere and built something of a client journey map or a matter roadmap, because clients need a clear picture of what working with you is going to be like, what's going to happen, in what order, what they should expect from each stage, including the time it takes. And I think setting those expectations can help them rest a little easier and maybe avoid the temptation of turning to AI to brain dump all their questions into.

And so I'll augment my advice from episode 108 and say that I think your client journey map or your matter roadmap needs a section about using your own AI for your legal case. And again, this is not necessarily a prohibition, although that might be where you start. It needs to be more of an orientation, right? This plain language rundown of exactly the six problems I just walked through written in terms your client can understand.

And frankly, I would write that document assuming that your client is going to feed it into their AI as well, because a lot of them will. That means making it structured, making it self-contained, don't assume too much background knowledge. And here's the thing, right? Approaching it in a way

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

that makes the document readable by a machine might just help make it more useful to a stressed out human at 11:00 at night.

But the document itself probably isn't enough, right? This also needs to be a conversation at your consultation, at a kickoff meeting, at an early check-in, wherever it fits naturally in your process. But the information is going to land better when it comes from you as part of a conversation, a human interaction, especially with a chance for them to ask questions and express concerns, than it is if you just lay it out in a CYA document.

And this is the essence of lawyering as caregiving, another topic I hit on a lot. And it's a way to make sure that the document becomes the artifact as opposed to the load-bearing component that you're not sure whether it ever got read or not.

Now, there are a couple of next level ideas I'm going to throw out there, and these are kind of just spitballing, but I think they have some merit. So I'm going to put them out there and let you think about them. One is to think about creating context documents that you've specifically designed to be uploaded into your client's AI. And these can be materials that steer how the machine responds about the questions or the information that the client is typing in and asking about their case. Context documents are kind of all the rage right now anyway, and they can be a really useful way to sort of set the table and set some boundaries for how an AI is going to approach a problem.

And I don't know, I'm really just spitballing here. Maybe providing those context documents strengthens an attorney-client communication or attorney work product claim, or maybe you're just setting them up for more discoverable stuff. I'm not sure. I'm not your ethics lawyer, and we're in frontier time, so it's hard to say what the right answer is, but it's probably something you can think through.

Another approach, and I've had some of my own consulting clients think about this, is you might be able to bring the client into your firm's AI tool, a

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

seat on your enterprise account, or even a dedicated machine in your office for them to use, something where the details of the interaction are somewhat managed rather than hoped for.

Basically, in that scenario, the message would be if you're going to use AI, at least use this safer one. And there's definitely some logistics to figure out. So obviously, it would have some convenience trade-offs for the client, but it might be worth talking with your IT team about, I don't know. Like I said, we're in these frontier times, and all of these things raise a lot of questions that really nobody has clean answers for yet. Records, supervision, what it means to endorse a tool.

But AI is the reality for a growing number of people, so we're going to have to learn how to work with it as accurately, as safely, and as humanly as possible, even if those things are sometimes in tension with each other.

So let me bring this back around to where I started with at least some hypothetical about Jason. And I think the client who's using their own AI on their case isn't necessarily going around you, they're doing exactly what you'd expect an engaged client to do. They're trying to better understand their own situation using the tools at their disposal. Jason almost certainly would have used these tools if he were in the same situation today.

So every question that he took to the law library, every statute he read himself, every hour of playing Matlock, all of that probably would have gone into a chat window instead. And you can't really stop that. And I'm not sure 100% that you should want to.

But what you can do is acknowledge the reality of the situation, educate and inform your client as best you can, and then create the best possible experience for them in their work with you. Really focus on building that human-to-human connection, because that's what's going to make the client less likely to supplant your judgment with the predictive text of an answer bot.

Ep #133: How to Talk to Clients About AI [AI ROI Part 4]

One last thing before I wrap up. As I was writing this episode, I fed these ideas into my own AI, in this case Claude, and asked it to help me create a one-pager for you to swipe from my website in case you wanted to use it in your own client roadmap documents.

Basically, it's a good first draft at a client-facing what to know about AI in your legal case insert that you can adopt for your own welcome packet, whatever your documentation is, plus some talking points for you to consider when you have the big talk about AI with your client. It's on the show notes page of my website, so you can go to agileattorney.com/133 to access it. And frankly, while you're there, you could also grab this episode's transcript and throw it into your own AI. Such is the modern world.

If today's episode gave you some good ideas about how to talk with your clients or think about your clients using their own AI, do me a favor and share it with a colleague. I see a lot of lawyers out there on LinkedIn, on Facebook, on Reddit who are really frustrated by their clients' use of AI. And I get the frustration, but I don't think it's particularly helpful. It's a big part of why I wanted to hit on this topic to begin with.

Even if you don't agree with everything I've said today, I think the conversation gets a lot better if more of us are having it together. And if you don't agree, please let me know why. You can email me at john.grant@greenline.legal. Of course, if you do agree, I'd love to hear that too, from a human, not a machine.

All right, that's it for today. As always, this podcast gets production support from the fantastic team at Digital Freedom Productions, and our theme music is "Hello" by Lunareh. Thanks for listening, and I will catch you again next time.