

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls



Full Episode Transcript

With Your Host

John E. Grant

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

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

I recently got a demo of an AI tool that promises to allow you to upload a legal document or a set of legal documents, and then the tool will automatically generate a series of tasks and build a workflow for you and your team to complete those tasks. It sounds like a productivity dream, right? But as a process improvement expert, I've got some really big reservations about how useful these tools will be in practice.

Based on my knowledge and experience, I can almost guarantee that these types of tools are going to come with unintended consequences that risk creating more problems than they actually solve. In today's episode, I'll tell you what those risks are and how you can mitigate them. I'll talk about the importance of understanding process as practiced instead of relying on process as designed. And I'll introduce the concept of task engines and why, despite maybe sounding appealing, they are actually something to be avoided whenever possible.

You're listening to the *Agile Attorney Podcast*, powered by Agile Attorney Consulting. 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, the other day, I got a demo of a legal AI tool, and I'm not going to talk about which one because I don't think it really matters.

And I'm going to be maybe not critical, but I think that there are some unintended consequences, or at least a high likelihood of some unintended consequences, stemming from what some of these legal AI and also general AI companies and technologies are trying to do, specifically around some of these newer developments in agentic AI and also using AI to define and fill up different workflows around, again, specifically the legal space.

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

And as you all know, right, I'm a fan of AI. I am a user of AI. I think that there are a lot of places where these particular tools, the LLM models, the agentic models, can really provide some benefit to lawyers, legal teams, and ultimately to access to justice and the availability of legal services more broadly.

But I think that we're making some assumptions about how AI is going to work, or at least the technology teams are making these assumptions, and they're then using those assumptions to sell lawyers and law firms and legal departments on these assumptions. And I think that there's some flawed or at least the reasoning behind the assumptions is missing some important context. So, let me unpack that for you.

The general premise behind most AI tools is that they will help you get chunks of work done faster. And they absolutely can do that. Research tasks, drafting tasks, all of them are really well suited for some of the capabilities of these large language model AI tools. The various pitfalls notwithstanding, right? We still have a hallucination problem and frankly, despite the claims of Damien Riehl back in episode 60 where he said that a lot of the hallucination things are being taken care of and he represents Vincent AI, which is just one of many tools.

I've actually had a whole series of LinkedIn posts and a discussion that was really popular on LinkedIn around how with ChatGPT specifically, I've actually observed the hallucination problem getting much, much worse. And it's wrapped up in the sycophancy problem, which is where AI tells you what you want to hear.

You know, there's lots and lots of issues still, and we're still working through them, and different tools have different ways of working around and mitigating those solutions, and frankly, they're all still really complicated. So I don't want to downplay, even again, despite the fact that there are some people involved who I respect and trust and believe are working hard on solving these problems, they're not quite solved yet.

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

That said, they can be mitigated, and I don't think that they should necessarily be a reason not to use these tools. I just want to make sure everyone is using them with a weather eye on the sky, a grain of salt, whatever metaphor works for you.

The other thing that I think these AI tools are really good at, and frankly, the best use of them for me so far is that they really help with writer's block, with the blank page problem of like, where do I begin on a thing?

And, you know, I'm using AI increasingly in various parts of my business to help me create a bad first draft so that then I can at least find a way in sort of mentally, psychologically, whatever, to engage with the content and work on it through an editing lens as opposed to an original drafting lens. And for me at least, and I know a lot of other people feel this way, that really helps me get into the flow of writing better.

Now, the reasons why it helps me get into the flow are interesting sometimes. I have often said that I am never as efficient as a writer as when I am hate editing something that AI has drafted.

So there's a lot about what AI spits out that I really don't like. It doesn't comport with my style. It, I think, winds up being really generic or jargony. And so, you know, I can like bash away at my keyboard when I am trying to edit something out that I don't like from an AI, but again, that's a pretty productive state for me. At least now I'm working on the thing and getting it to say what I really want to say. So that bad first draft can be useful.

All of that said, these are all examples of using AI to speed up your working time on particular projects or deliverables, right? So you can use AI to become a faster researcher, to become a faster drafter, to become a faster editor, and all of that is great, and it does in fact work for that, although maybe not quite as well as originally sold. And I think it can vary based on different use cases.

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

But when I'm looking at process and deliverables as a process improvement person, I'm not just trying to speed up the working time for a particular piece of work. I want to improve the overall cycle time, right? So the amount of time it takes from start to finish to actually deliver the work to its intended destination, whether that's to a client, to a court, to an opposing counsel, to an internal resource, whatever it happens to be.

The working time is just one part of the equation in the total time, and the other part of the equation is waiting time. How long does the work spend inside of your system waiting for a resource to become available to work on it or to do quality review on it?

I talked about this back in episode 40 where I discussed my metaphor that I often urge my clients to adopt, which is imagine a GoPro camera sitting on the work itself, on the deliverable. And anytime the client can log in and see what's happening on their matter, and you know, 80, 90, 95% of the time, what they're going to see is the work is just sitting there waiting for something to happen.

And so when I'm looking at process improvement, a lot of the time that I spend is figuring out how can we reduce the wait times, not how can we simplify or speed up the working times. Now, there's a relationship there, right? One of the reasons that the work is waiting is that the resources are working on other things.

So if you can speed up working times, then yes, in theory, you can also then start to reduce the waiting times, except we have this problem, and that problem comes from Little's law, which I've talked about a lot on this podcast. I've referred to its cousin, the Kingman's formula.

But just as a review, basically what Little's law tells us is that the rate at which you can deliver a piece of work in your system is proportional to the total amount of work inside of that system. So said another way, the more

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

work you have in progress inside of your overall workflow, the longer it's going to take you to deliver any one piece of work.

And I talked about this back in episode 73. My concern with technology tools in general, but this AI boom and sort of gold rush right now, is that if we think that the AI is going to help us get more done and therefore we let more work into our system, then the counterintuitive thing is going to happen, which is even though we might speed up the working time for the deliverables itself, we will still take longer overall to deliver the work because we're letting so much work into our workflow, into our finite capacity.

So that takes me back to this demo I got the other day where the person showing me this particular tool was really excited about this idea that you could take a legal document or a set of legal documents, and they could be pleadings, they could be contracts for review, they could be demand letters, whatever it happens to be, and you can upload them into this AI tool and not only will the tool sort of summarize what's going on in that particular document, it will then suggest a series of tasks that you should or your team should be undertaking in order to respond to that document.

Now, it's that generation of tasks that has me of two minds, right? On the one hand, I loves me a good action plan, right? I really think that being able to break down a series of work items and put them in a logical order, I think that is great. I think it's important. I think it's something that lawyers frankly don't get a lot of training on. These are sort of fundamental tools of project management and process improvement, and it's stuff that we all could stand to be better at.

That said, generating tasks feels a lot to me like generating work and creating more work in progress or potential work or should-dos in your law practice. And if we aren't doing that in sort of an intentional and a vetted way, then I'm worried that we've just sort of created the conditions for Little's law to kick in and really slow the overall progress of work down even

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

though the speed at which we can accomplish some of the tasks has gotten faster.

And I'll flip over here for a minute and talk about some recent headlines from a different AI tool. And I'll name this one because I've talked about it already in other places. But Harvey AI, which is a big law product. It's owned by or developed by one of the big law firms, I think. I don't pay as close attention to the big law tech as I do the kind of things that are impacting most everyday lawyers, people law lawyers.

But the headline from this Harvey press release was that Harvey is going to be able to help legal teams define their workflows and then through the use of agentic AI, right, which is basically saying, okay, that the software is acting as my agent and is going out kind of into the world of the internet, not the broader world, but that's blurry lines these days. But so the AI is going to go out into the world and accomplish a series of tasks for you, right?

So when a certain event happens, it will go into your maybe email program, draft a client communication that talks about how you just finished reviewing this contract and here's what you found, and then it's going to send that update email to your client automatically. And then that's going to take a series of tasks off of your plate and maybe solve the client communication problem that so many lawyers and law firms have.

Maybe. I talked about my concerns around that type of thing back in episode 63, which the title is Legal Automation Pitfalls. And I have my doubts about the extent to which an auto-generated email or other sort of client communications are going to actually solve the client needs around client communication.

And my concern is again, that the things that are delivered by the AI, even this agentic AI in an automated way, are going to create task engines that

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

kick off even more or unexpected things for the lawyer or the team to do that wind up disrupting workflows instead of streamlining them.

And even that, right, is a solvable problem. It's something that with a few different iterations of using these automations, of crafting these communications, the tools can get better and better at it. But it's going to take more effort and more work and more iterations than most people think it's going to.

Now, my other concern around this announcement from Harvey, and again, Harvey's far from the only one, there are lots of AI companies and technology companies that are purporting to be able to use their tools to not just improve your workflows, but to actually define them for you.

And there's a very real problem in the world of workflow design and process design and process improvement that is as follows. It's the gap between process as designed and process as actually practiced. And I'll give you sort of the classic example of what we mean by that.

And it's the pathways that cross a college campus, right? A quad or another sort of open green space. It doesn't have to be a college campus. It could be a corporate campus. But if you imagine sort of a bird's eye view of these pathways crisscrossing a green field or other setup, and the idea is that the people are going to use these paths to get from one place to another along the most efficient route.

But if you look on most campuses, you will find that despite the best intentions of the architects, of the designers, everybody else involved in the planning process, there are almost always dirt paths that connect different parts of the campus that were not part of the original design.

So when I go into an organization and I'm about to do process improvement work with them, one of the tenets of the Kanban method and other methodologies is to start with what is done now, start with what you are doing now.

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

And oftentimes the law firm leaders, the managers will say, oh, well, I'll just send you our current SOPs and you can see what we're doing now. And I do not accept that as gospel, right? It's important, it's information, it's useful, but I need to actually observe inside of the team what they're really doing, not what the policy and procedure says they should be doing. And there's two reasons for that.

The first one, and the most likely one frankly, is the people doing the work have often discovered better ways of doing the work than the policy or the process envisioned. It's really hard to sit in a room with a whiteboard or stare at a blank space on a keyboard or on a computer screen and come up with the ideal process as actually practiced, right?

We all have an optimism bias, we all make assumptions when we're doing this design work, and it's really hard to know what is actually going to work or how things are actually going to go until you put it into practice.

The other thing that I'm trying to uncover, and this is really common as well, is that oftentimes the senior attorneys or the managers, whoever is designing the process, will make a lot of assumptions about what the team members know or should know about what needs to be happening. And they will mistakenly believe that just because they've created a policy around what needs to happen, that the members of the team actually understand the reasoning behind why that needs to happen.

And, you know, it's my experience in doing process improvement with teams in particular that it is far more effective in the long term to make sure that the people doing the work understand why things are important, right? What are the reasons behind the things we are doing, than it is that they understand a point by point recipe or a process action step for what they're supposed to be doing, right?

I'd rather people turn their brains on and understand the reasoning and then use their creativity to come up with good ways to accomplish the goals

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

of that workflow or process, then to have them turn their brains off and just follow a step by step plan for how some quote unquote expert believes that the work should be done.

And this is my concern with tools like Harvey or other AI that purport to be able to define your workflow is I just, number one, I don't think that they are in a position yet to be able to really understand the work as it's actually done right now as opposed to the work as designed. Number two, I worry that these tools are focused on pushing out these process steps and they're really overly focused on the step-by-step recipe without doing a good job of communicating the reasoning, the why behind the work.

Okay, so back to this other tool, the one I got the demo for that is sort of creating the task steps, which are related to the workflow steps. And I'm going to use a word that I used in passing a minute ago.

But I really want you to think about the concept of task engines. And this is another term that was coined by Cal Newport. It was in his book, *Slow Productivity*, which you all know I am a big fan of that book. If you're going to read a book this summer, that would be the one I recommend. Even though just last week I said you should read fiction and other things, not workbooks, and so do that too, but if you are inclined to read a book about work, read *Slow Productivity* if you haven't already.

But when Cal Newport talks about task engines, he's talking about projects that generate a large number of recurring tasks, right, requests, questions, meetings, little deliverables that will eat up your time and attention but not always make clear progress towards a desired result.

And this goes back to some of the core tenets that Newport and I share. Right, number one, what I refer to as the honest reckoning with capacity. We only have so much time, attention, energy to devote to all of the things that we want to try to accomplish in our day-to-day, week-to-week, month-to-month, year-to-year lives.

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

Number two, that capacity is probably smaller than you think, right? I keep talking about the optimism bias, and we all want to be able to do and we want to create a world where we can quote unquote get more done or accomplish more things than we probably realistically are able to accomplish.

And the reason that we need to have the honest reckoning with capacity, ideally using feedback loops like the daily standup, the weekly planning meetings, these weekly reviews, the periodic retrospectives, is to recalibrate our actual sense of what we can and can't accomplish in a given time period. And then the whole point behind agile methodologies is to shrink those time periods as much as possible because we're better at estimating our capacity in small chunks than we are in large chunks.

Number three, once we have that honest reckoning with capacity, we are more realistic, hopefully a little bit data driven around what that capacity really is, then we have to be more brutal about our assessment of our priorities. We have to be more intentional about what projects, what work we will commit to and what work we will say no to or we will defer to some later time.

And so the caution that Cal Newport gives in *Slow Productivity* to avoid task engines is when you're doing that assessment of priorities, when you're deciding what work you're going to commit to, you need to consider how much work the overall project is going to generate for you.

And his core idea is that you should be clear about wanting to accept the kinds of projects, and when I say accept, these may be things you're delivering or you're creating for yourselves. It's not necessarily things that other people are pushing on you.

But ideally, you want to be able to select projects that are going to let you accomplish larger goals with a smaller task footprint, right? Fewer actual

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

little interstitial things that you need to do that will suck up your time and attention without maybe as much return on that investment.

And so when I saw this demo of this AI tool that is spitting out a bunch of tasks, my first thought is, "Oh my gosh, that's a task engine." Now, that's not necessarily a bad thing, right? There are client work or deliverables, projects, matters, whatever that frankly do generate a lot of tasks. So I'm not saying that we can avoid all of those, especially if it is literally our job to do them.

But I've kind of seen this movie before and, you know, I'm going to call out a software company that I am generally a fan of and again, I like the people, I like the product for the most part, but Clio is a task engine. And particularly the way that Clio conceives and creates tasks, the more you sort of invest in setting up the various task templates and the conditional deadline reporting stuff that it does or the creation of deadlines, all of that is great. It's automation, it's not AI based automation, it's just logic trees, but it's the same kind of basic idea.

Clio and other tools, it's not just Clio, are really good at generating a certain quantity of task items for people on your team to do and Clio would say better yet, I would say maybe worse yet, it is also really good at assigning those tasks to people on your team and setting the deadlines for those tasks.

Now, again, this is like at the root of what we think effective lawyering is, right? We certainly want to make sure that we understand what deadlines exist. We want to make sure that those are calendared, right? There's any number of treatises around avoiding malpractice risk that basically boil down to make sure you're calendaring your deadlines really, really well.

So I'm not saying that's wrong, but here's the problem with an automated task generation tool like Clio or like this AI tool is that if all it's doing is generating tasks and even setting deadlines for those tasks without

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

accounting for the commitments that your team has already made and the finite capacity that it has, then you actually are running the risk of overburdening your team or certainly overburdening them relative to certain deadlines and certain deliverable dates.

And this is the problem with Clio is it has no concept of capacity. It doesn't care if your paralegal already has 87 tasks on their plate, it will happily add 14 more based on this automation that just got triggered. And I believe that is true of this AI tool as well. What's worse is I'm not sure that this AI tool can go so far as to actually understand who is going to accomplish certain tasks. And again, over time, it will be able to do that. Some of this is going to be just the iterative development of the programs, but they are not there yet.

So I'm going to leave it there for now. But my overall theme for this episode and believe it or not, this is actually a continuation of last week's, right, where I encouraged you to slow down, embrace the seasonality of summer, use that time to be intentional about what you reasonably hope to accomplish with your law practice in the second half of the year.

And then I also talked about doing that as a way of inoculating yourself against the coming barrage of marketing messages and hype around legal tech that you know is, it 's happening now, right? I'm not pretending that it isn't happening over the summer, but I think we all know that it's going to accelerate come September.

But I'm hoping this episode will give you some tools to help you sort of better assess the claims and maybe see through some of the hype surrounding the claims that these technology companies are about to make.

Number one, getting parts of the work done faster is not the same thing as improving the throughput or the cycle time or the delivery time of the work overall. Just because you can get certain tasks done faster does not mean

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

that your overall practice has become more efficient. It only matters if the end-to-end deliverables are going faster.

Number two, speeding up your working time is not necessarily going to improve your cycle time if you can't also reduce the waiting time, the amount of time that work is in queue waiting for something to happen as opposed to actually getting worked on.

And the only way that I know reliably to reduce waiting times is to put less work in process at once. And by accepting fewer matters or fewer projects, fewer tasks overall, you're more likely to get a handful of tasks done faster, and that will counterintuitively allow you to speed up your overall delivery rate because you're focused on getting things done as opposed to getting things started.

And the specific place where this is going to come into play with AI is the extent to which people need to do quality assurance review on AI-generated work. And I've talked about time and again, the two places where I see bottlenecks form the most inside of law practices, number one is client homework and number two is quality assurance review. So it's a very real risk.

The last thing is to beware of these AI-generated or technology-generated, it doesn't have to be AI, but these technology tools that are generating workflows for you or task sets for you.

And again, I think it can be great for helping solve the blank page problem. I don't know how to document my workflow, and so I'm going to spit a bunch of things into an AI tool and let it give me a first pass of what it thinks the workflow is that I just described. That can be really useful, but don't end there.

You have to use that as your starting point and then really interrogate, okay, well, this is what might be happening, but how do I figure out what's really happening? And the only way to do that is to observe. You can talk to the

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

people on your team, but they're not always reliable narrators, right? The best way to do it is to observe.

And I'll put in a plug for a high-quality Kanban tool, which as I've talked about before is different from just using a Kanban interface, but a good quality Kanban tool will use the metrics and allow you to analyze how work is flowing through your particular workflow as defined by your Kanban board and that can be a really useful tool for reality checking what's really happening inside of that workflow as opposed to your assumptions about what you believe should be happening about that workflow.

And I guess the final, final thing is beware of task engines, right? When your capacity is finite, you want to make sure that you are organizing your team and your work and also accepting work into that finite capacity, keeping in mind what needs to happen to actually move it forward.

And the kinds of tools and processes and people, frankly, that are task engines that like to generate a high volume of little checkpoints, little deliverables, whatever it is that takes up your time and attention, that's not always going to result, in fact, it's often not going to result in the most efficient way of getting things done.

To put it another way, even tools that help promote worker efficiency, right, making sure that the people on your team are delivering work better isn't the same thing as promoting flow efficiency, which is the end-to-end deliverable of work through your system.

All right, that's it for today. As always, if you have questions about this, if you want to talk about how to apply some of these things in your own law practice, if you want to sit down and do some workflow analysis or come up with some strategies for how to better adopt technology in your own law practice, you're welcome to set up a discovery call with me. You can find information about that on my website, agileattorney.com, or you can just shoot me an email at john.grant@agileattorney.com.

Ep #80: AI for Legal Workflows: The Potential and the Pitfalls

As always, this podcast gets production support from the amazing team at Digital Freedom Productions, and our theme music is "Hello" by Lunara. Thanks for listening, and I will catch you next week.