

**Artificial Intelligence and the Practice of Law**  
**LawCast BC podcast**  
**Law Society of British Columbia**

---

**Vinnie Yuen:**

Welcome to LawCast BC, a podcast produced by the Law Society of British Columbia. The Law Society regulates lawyers in BC. Our mandate is to protect the public. I'm Vinnie Yuen, your host and producer. We're discussing a very interesting topic today, artificial intelligence and the practice of law.

The concept of artificial intelligence is not new. The academic discipline of artificial intelligence was first established at Dartmouth College in 1956. Of course we've seen numerous advancements since then. Today, there are thousands of AI tools including generative AI tools like ChatGPT and Google's Gemini. So how do these AI tools fit into the practice of law? Today, I've invited Jon Festinger, KC, and Robert Diab to come chat about the ways that AI has been used to help practice law, the risks in using AI tools, and how AI could potentially change how people access legal services. Jon and Robert are working together to put together a course on AI, law and justice which will be taught by them next year at Thompson Rivers University. Jon Festinger, KC, is a Vancouver based counsel and educator as an adjunct professor at UBC's Allard School of Law. He has taught a wide variety of law courses relating to intellectual property, media and communications and business for more than 30 years. He also teaches at Thompson Rivers University. Jon practices law as counsel at the law firm of Chandler, Fogden, Lyman. You may also remember Jon as the previous host of our Rule of Law Matters podcast. Robert Diab is a professor at Thompson Rivers University's Faculty of Law. He writes about constitutional and human rights and topics on law and technology. This includes work on privacy, encryption and AI and on powers of detention, search and public order policing. Prior to teaching at TRU, Robert practiced criminal and administrative law in Vancouver. Here's our chat.

Thank you so much for being here Jon and Robert. Can you please tell us a little bit about yourselves, how long you've been writing about or teaching students on the topic of artificial intelligence?

**Jon Festinger:**

I've been teaching on the subject of media and the law and law schools, specifically at what is now the Allard School of Law, since 1993. I guess that AI really started coming into my courses in a very rudimentary form when I started teaching videogame law in 2005 because early forms of artificial intelligence were embedded in videogames since about the mid-1990s. So that's when the, kind of the first wave of speculation around AI, and then sort of I'd say a second wave, which was very much sort of data and algorithm driven, and I'm sure we'll get into where the technologies have taken us today, I was certainly speaking and teaching a lot about it in say 2015, '16 onwards. And then the third wave of AI teaching and writing I would say really came along with ChatGPT 3.5 and then 4 where everybody and every profession wanted to talk about and understand practical implications as opposed to just you know the theoretical issues around AI which is what we were, I think we were all in academia concerning ourselves with initially. Over to Robert.

**Robert Diab:**

Thanks Vinnie. I have been teaching law at Thompson Rivers University since 2012 and I have not taught law and AI yet, I'm planning to with Jon next year, so my interest in it is, has mainly been you know on the scholarship side. I've started writing about it in a few places. I'm following it you know closely, I'm reading everything I can about it, I've very curious about how it's going to shape the practice and the teaching of law. I think it's already having an impact on the way students are learning the law. I know sort of from first hand accounts that a number of our students are making use of it in various ways and I'm following how practitioners are making use of it. And so I think it's just something that we, you know we have to grapple with and to try to be on top of as soon as we can.

I think that generative AI marks a clear break in the development of you could say the digital revolution. I mean I think this is really on the scale of something like the advent of the browser or the web, you know it is a, it is a change of that magnitude you know.

**Jon Festinger:**

Can I just add something to what Robert said, and that is where do we place AI sort of in the history of human development of tools? Maybe I'm too radical on this and maybe it's a little counterintuitive, you know as somebody who's you know practiced and taught communications law for kind of most of my career, I look at the web as actually more or less a natural progression that started with the telegraph. When I look at AI, I was thinking about it today, you know where do we put it in the, in the firmament of sort of human development and I think it would overstate it to say it was important as the invention of fire for example, but I think it belongs with the invention of the clock for example, in a sense the invention of time as we know it, mechanical time, the ability to measure time you know has completely transformed our world. For me, and maybe I'm overstating it, but that's where I put, that's where I put artificial intelligence right now.

**Vinnie Yuen:**

Given that concept, given how much weight you've given to AI tools, how have we been seeing that play out in the practice of law? How have we been seeing lawyers use AI to help them practice law?

**Robert Diab:**

So my sense is that it's early days, it's, we're very much still waiting I think for the tools to become better tailored to law in particular, you know and to the law in a given jurisdiction and to the tools to become more accessible to us you know in our, in our place where we practice. And the tools that are available I think are rapidly evolving so just to give a general sense, I mean you know we still, we haven't yet seen a language model be connected to CanLII, you know the sort of the main source that lawyers use for legal research. We have yet to really see language models be pervasively connected to the really general tools like Word, you know Office, that's coming but it's not yet here.

There are a handful of companies around the world that are trying to tailor this technology to law in particular to make it hallucinate less, to make it you know conduct legal research effectively, to maybe produce briefs, to you know read a statement of claim and produce a reply or you know to produce a statement of claim in the first place given a fact pattern, that sort of thing. So what we're only in the process now of seeing you know the first batch of companies create these tools and you know they're being rolled out selectively, they're in beta, etcetera, so all that's you know still unfolding.

But we are getting a sense, you know we can see by using more generally available tools like you know sites like Perplexity AI or Bing or sites that allow you to upload a PDF which could include a case, for example, and have it summarize the case or you know the brief, etcetera, for you. We're just getting a taste of how powerful these tools can be but I think the open question that remains at the moment is will they succeed or will, you know will one or another company succeed in making these tools effective, reliable, really useful to lawyers. This is a point we can come back to but I think we're still, we're still waiting and it's not entirely clear how they're going to be used to help us in practice and you know to what extent.

**Jon Festinger:**

What Robert said I think was particularly noteworthy is we're still waiting to see how powerful these tools can be. We're not there. So to answer your question directly Vinnie, some lawyers are using the tools well and some lawyers are using the tools not so well. Using the tools well, in my view, means a couple of things. It means using the tools to prompt information or ideas that you might not otherwise have to think horizontally about a subject, to prompt maybe analytic thoughts that would be helpful to your client or helpful to the rule of law. That's using it well.

The second part of using it well, for reasons that Robert has well explained, is check absolutely everything. Because of hallucinations, you really don't know what you're getting. In some ways, and this isn't the first tool that's supposedly gonna save lawyers

time that could add time to what we're doing because you literally have to check everything because citations can be made up, cases can be made up, perspectives can be flat out wrong you know and I would commend any lawyer to ask a generative AI to do something in their area of greatest expertise and then see how it does because that's, then you'll get a real sense, if you're asking it about things outside your area of expertise, of how well or badly it's doing.

So that's we use it well. How you use it badly is essentially the inverse of that and we've seen this in various jurisdictions including British Columbia where lawyers have perhaps tried to use AI as if it were Google, as if it were going to, in a linear way, give you something that was at least factually correct or that you could have confidence in being factually correct whereas that's not what AI is. You know think of AI, I would say at least at the current stage, as an irresponsible 14-year-old that you're asking questions to and will sometimes tell you what it thinks you want to hear and sometimes will just make stuff up and sometimes will run away from you and sometimes will run towards you. That's kind of the level we're at and are we going to get beyond that where we have a really useful tool that helps us think laterally. I think we will as long as we're looking at it as thinking laterally. If we're looking at AI and saying write me a brief, write me a factum, I don't know that, this is my view and Robert can disagree, I'm not sure that we're ever gonna be in a place if we don't bring our human thought to the particular facts and the particular case and the particular client we have. I don't know that we'll ever be in a place where I would feel comfortable saying hey, you know the AI did this, I don't need to check it, I know that it's better than what I can do. Robert, what do you think?

**Robert Diab:**

Well just maybe two followup thoughts perhaps to come at this from a different angle, maybe it might help to kind of make it a bit more concrete. So just in terms of where AI is at the moment in the practice of law, as far as I'm aware there is, there are a few, there's a few language models made available to lawyers in Canada, one is Lexis, the Lexis tool which is proprietary. There's a fascinating new development out of Queens, there's an

open source legal language model that's available at the moment to law students. It's called Open Justice, if you Google that.

I was playing with it over the last couple of days and I ran a few searches. So I was using it, to be clear, I was using it as a research tool, I was using it as an alternative to going onto CanLII and running search strings to get an answer. And one question I asked it, so just to give you, to make this concrete more specific, I asked it what are the factors a court will consider to decide whether you were in care and control in an impaired driving prosecution and give me the leading case. And it, now that, to get to that prompt, I had to kind of run two or three queries first because I, it sent me off in a wrong direction, it gave me a summary of general negligence law. But anyway, once I did narrow the query down, it gave me this remarkable like 400 word summary of the leading case *R v Boudreau* from the Supreme Court of Canada 2012, which it didn't cite like you know McGill style but it gave me the case, the year, the court, and it gave me a very good summary of the rule and the factors.

But then I ran another query, something about give me you know a summary of the cases where the court used, where the court considers whether the power to search incident to detention can be used when a suspect throws a gun away, like a really, really specific fact pattern. So in other words, spare me having to read 30 cases on CanLII, give it to me in a few words. And it was, it was so-so, it gave me a couple of cases but the rest of it was not help...like so let's just say about, there was a, there was 40 percent of useful stuff out of you know the whole answer. So that's just an example of how you know we're getting a glimpse of the promise of these tools.

The question though, at the end of the exercise, was would I have been better off just going to CanLII and running a search string, like did I save any time, and maybe more to the point, if I didn't know the area well enough to be able to prompt, to be able to craft the prompt effectively then you know was this, was this tool really useful. Like in a lot of

ways, you get the sense when you use a chat bot that the, your question contains the answer and I think that's the conundrum with law, your effectiveness with these tools is really going to depend on how well you already have internalized the area of law you're working with.

But this, so this leads to the second point I wanted to make in response to Jon and that is that there are many different tools here, AI is not one tool, you know it won't be the case that in five years, if the practice of law is considerably transformed by AI, you're gonna have one app that does everything for you, you're gonna have a Swiss army knife app. You might use you know, you might use 10 different tools, one tool to help you clean up your factum, clean up the writing, maybe another tool to do some research, maybe another tool to give you you know ideas for you know for how to craft a document, produce a contract, etcetera. So it's all still very much unclear but I, but I doubt it will, it will be one app or one solution.

**Jon Festinger:**

I think we're very much on the same page. I do want to add that the example you gave really goes to what, the point that I would like to make is that the application of the law to a particular set of facts, which is what we do with lawyers, is the application of law to a very particular set of facts that for a very long time, and certainly my lifetime and I hope you know you and Vinnie live very long lives, but it might include your lifetime before we get to the point where we have an AI that will reliably allow for the law to be applied to a particular set of facts correctly enough so that it is reliable, so that lawyers don't have to second guess it. Where I think AI actually shines even today is in terms of giving us lateral ideas, you know prompting things in us, we always hear about us prompting AI but AI prompting us in its kind of, in what it returns to us, to think about things that maybe weren't within our field of view previously that could be beneficial to a client. But to do that well, you already have to be a pretty good lawyer, there is kind of a threshold here, you have to be a good, experienced lawyer to be able to do that but AI does that today. AI, one thing that I've learned in some of the experiments that we've done here at Allard is AI asks questions really, really well.

**Vinnie Yuen:**

So we talked about the risks of using AI, we talked about how you know we need to check the results, we recently received notice from the Federal Court that requires parties to declare if their documents were prepared with the help of generative AI; what do you think about this development? Is it a step in the right direction?

**Robert Diab:**

Yeah, that was a fascinating development. So for you know for listeners who may not have noticed this, the Federal Court is just one among many courts that have issued a directive recently about the use of AI in litigation. And in the case of the Federal Court, it's done something interesting, it's said that if you use AI in the preparation of a document, well it's a certain list of documents, not every document but just certain documents, like for example, a pleading or a factum or like a memorandum of all arguments say, and it also carefully defines what constitutes AI, you know like spell check wouldn't be AI but having a language model produce you know two to three paragraphs of your statement of claim would be AI. And what it said is that it wants a declaration that you have used it in the preparation of X so, and it gives an example, it says you know if paragraphs 25 to 30 of our statement of claim were produced by AI, I am not sure that this is the best approach.

It is, it stands in contrast to the approach of other jurisdictions where the directive is if you've used AI you know you must assume, you must confirm that everything said is correct or that you stand behind what it is. So I, I find the latter approach more appealing. I'm not entirely sure what the point of a declaration of the fact that you've used AI in paragraphs X to Y is other than to essentially call into question the, you know the merits of what you're doing. And the odd thing is that the, the directive goes on to say, and there's not that there's anything wrong with using AI or that using AI you know gives rise to a presumption that it is in some way flawed, your document, your argument, whatever. So I think the best approach is to say what a number of Law Societies have said, simply look, you are responsible, so you may use these tools as you would a spell check or whatever else, a word processor, but it's your work and so you will be held responsible.



**Jon Festinger:**

I'm going to take that into a slightly different realm where this question comes up, and Robert it's your and my realm in, in a different way, and that is in education, how do we deal with students submitting AI papers? You know we have a very disconcerting piece of research that just came out, and it seems like very reliable research, that shows that academics with our egos, we think we can tell if AI is being used by a student in the generation of a paper and that's now been studied and we can't. And what I've always told my students is AI is a source, treat it as a source so you have to footnote it. Very few students do which leads to the question of you know are they, are they using it and not disclosing because you know they fear that they will be judged more harshly? And that, the Federal Court rule does, as Robert points out, seem to raise that almost subconsciously, you must declare, you must effectively sign an affidavit that you are disclosing and then on the other hand, goes to some length to say but we're not judging you on that. It does seem odd, does that mean, and it sets an odd precedent, does it mean that every, because there are other things that can be wrong in legal research. Headnotes are sometimes wrong, legal textbooks can be wrong and incomplete, I hate to say that having written one but it can happen. And lawyers don't have to certify those and the purpose in law is to get it right in the end.

So you know we've talked about lawyers having to check everything, well judges and clerks have to check everything as well and that's part of the process. I agree with Robert's misgivings. I do think it is common answer in today's world to say we need disclosure and I do think disclosure is a very good idea so I do understand where the Federal Court is coming to its principle. Whether it leads to a better result for litigants, which after all, and for society, which is after all the purpose here, whether it adds to the process and the safe guards in the process or in some way that we don't see distorts the process, remains to be seen.

**Vinnie Yuen:**

There's obviously a lot of risk in using AI and are there examples where there have been liability concerns where the use of AI has landed some companies in legal trouble?

**Jon Festinger:**

Well the one liability case, the only actual bit of litigation or threatened litigation I know of is in the United States where a client who was a rapper, who is a rapper, was suing his lawyer for having lost a case on the basis of the lawyer used AI in the case but that's the only liability case around AI involving a lawyer. There is a wonderful court decision involving Air Canada where Air Canada's chat bot gave some, some advice, the chat bot being AI, to a, an Air Canada customer about an Air Canada policy and the chat bot was flat out wrong on the policy. When, when the customer went to Air Canada and said you know give me my compensation, Air Canada said that's not our policy, the customer said but your chat bot told me it was, Air Canada said but it isn't, the customer sued Air Canada and quite correctly, in my view, won where the court said Air Canada, you cannot disavow the mistake of your own chat bot.

**Robert Diab:**

Maybe the only thing to add to that is that I think we can already foresee a point in time when AI becomes good enough that lawyers are integrating into their practice to a significant degree and begin to rely upon it to some measure. And then I think at some point we're gonna start to see cases or disciplinary proceedings where the question is whether a lawyer who gave advice largely based on the product of AI who did not confirm it or did not you know sufficiently vet it, did the breach the standard of care, you know did they provide negligent advice in doing this. So I think our conception of what the standard of care requires will probably change over time but right now you know it's clearly you cannot rely on this in any way except to sort of you know start you off in research, give you ideas, help you you know hone your, your written materials but it isn't in any sense something can rely on.

**Jon Festinger:**

Can I just add one other thing about what makes AI allow us to feel like it's reliable when it isn't? AI is, because AI is coming from a survey of the web, you know if you say to an AI write me a factum for the BC Court of Appeal, it will purport to write you a factum for the BC Court of Appeal. It will probably be, as Robert recounted in talking about his experiments, it'll probably be 40 percent right and 60 percent wrong but it'll be 95 or even

98 percent correct on form, it'll look great. And we, as a society generally, but especially as lawyers, we sometimes value form over substance. Form's terribly important to us and because AI can generate form that is very good, very accurate, AI does this really, really well and as professionals, we really have to watch that.

**Vinnie Yuen:**

We've talked a lot about how lawyers can use AI tools; I'm just gonna change gears a little bit and talk about how the public, people who need legal services, whether AI can potentially change the way they access legal services. Will it improve access to justice? Is it a useful tool for the public?

**Robert Diab:**

I'm optimistic about this. I remember in advance of the, my final exams for the spring, I took an exam question that I had, like I had a short fact pattern and I plugged it into GPT because I wanted to see what the quality of the answers was partly out of curiosity but also because I had to decide you know should the exam writers have access to the open web and if so you know what would happen if they plugged this, this into GPT. And I remember being struck by the quality of the answer. So on the one hand, I remember the observation that for the layman, the answer that ChatGPT gave, even though it was you know I wanted like the answer in Canadian law, the answer was pretty much bang on. It was like if you were just Joe Citizen charged with this offense and this was the fact pattern, the answer was generally correct that this was, you know this was what would happen, this was the law that would apply. But for law students and lawyers, the answer was not helpful in the sense that it was only, it was too general, it gave you no specifics, it gave you no you know cases, provisions, etcetera.

I think that probably in the near future when language models are trained on more specific bodies of legal material and do become accessible in the way that for example OpenJustice is accessible to law students, I think that they will provide something useful to citizens who just want you know a general answer, a ballpark answer or maybe you know basic pointers about how litigation would unfold or where they should go to

advance their claim. I think it can provide them something useful in the way that for example, when you go to Perplexity AI and you run a query, pretty consistently it gives you the sense that instead of running a basic Google search and spending say 20 minutes reading the you know the answer in the top five sites, it's giving it to you there in 2 to 300 words. Like I think, I think we're gonna have legal tools very soon that will do that reliably for members of the public.

**Jon Festinger:**

To add to that and maybe give a little bit of perspective on that, everything that Robert said is a 110 percent correct but let's have a perspective on what that means. My starting point is that we've not done nearly well enough as a profession in empowering access to justice and creating true forms and effective forms of access to justice. So if we've not done very much in the real world, AI is fantastic because it adds a whole layer of knowledge and information and a toolkit, you know write me a statement of claim, write me a factum. You'll get something out of it which was not available to an impoverished litigant or a lay litigant before. But it's not gonna nearly, for all the reasons we've discussed so far, it's not nearly gonna be as good as it would be if done by a lawyer or done by a legal clinic on the current state of technology.

In some ways, you know what we're really gonna have to measure, and I don't have an answer to, is the gulf becoming bigger or staying the same because everybody now has the tools, or is this going to even things out a little bit more. I am optimistic like Robert that over time it'll even things out some more. I do think even on today's, in today's state of development, AI can be a step forward if we develop a purpose built AI for lay litigants. That's what we really need to do, we need to build an AI access to justice tool and that AI access to justice tool, I'm sorry to say, because of the way law is done in the world, is gonna have to be done jurisdiction by jurisdiction. I'm hopeful that somebody will come up with a great design that then can be taken into every individual jurisdiction and populated with information that works. But that's a real process, it's gonna take a couple of years even if we started today, and we should start today, to test it, to do it properly.

AI will, is already gonna take access to justice even in its current very rudimentary and incomplete form, it will take lay litigants to a better place than they were before but make no mistake of how far we really have to go to empower people in ways that don't put them at a significant disadvantage in front of the courts. We need to start with residential tenancy and employment issues. Robert and I have discussed this, we are a long way away from an AI that's gonna help you in an administrative law matter and unfortunately that also includes, because immigration law is largely administrative law, it's gonna be a long time because of all the factors and because of all the fairness factors. AI doesn't do fairness, at least not right now, I think we have to do a lot of work to figure out if we can get AI to do fairness. But employment law, residential tenancy law, tax law, I think there are areas where AI can help a fair bit fairly quickly.

**Vinnie Yuen:**

What aspects of legal services do you think could potentially be replaced with AI tools and then what services could never be replaced?

**Robert Diab:**

I agree with Jon that there are gonna be some areas of law where AI, both as a research tool and as a predictive tool in terms of outcomes, fits more easily than in other areas. But in a broader sense, I wonder whether the word replacement is as helpful as the idea of you know of a tool or an assistant so, or to put it in different language, other people speak of how it may be more helpful to think not of AI as an autopilot but as a co-pilot. I, at the moment, I have a hard time imagining that any area of the law will see what humans did being completely replaced by AI. I mean certainly you know the drudgery of going through like a million documents to look for X maybe that can be done entirely by AI, but in any meaningful sense, anything that involves judgement, you know insight, assessment, there's going to have to be a human in the loop and so I don't think we should fear that we're going to be rendered obsolete anytime soon.

**Jon Festinger:**

I totally agree with that, that there, the notion that lawyers are going to be obsolete and especially the notion that judges are gonna be obsolete and that we're gonna ever get to AI judging, I, you know I fear the day of a judge being an AI because that is a day where I literally fear for democracy, I literally fear for the rule of law.

**Vinnie Yuen:**

On that note, do you think that the future of the legal profession will look a lot more different because of AI or other technologies?

**Jon Festinger:**

So I will, as the, as among the two of us the person here who's been called to the bar for over 40 years say that, and I might be wrong but you know I remember when word processors came in and word processors were going to change everything and we lawyers were going to be, we were gonna have three day work weeks because there wasn't gonna be enough work for us and it did exactly the opposite. It just, it just made us bicycle harder and harder and harder. We were on the treadmill, you know we went from 1400 billable hours to 1800 billable hours to 2200 billable hours with every technological advancement because other things change. And you know AI, in some ways, will make the law more accessible to everyone but AI making the law more accessible to everyone also makes the need for lawyers greater.

You know when you look at the amount of medical information that's out there, are we really our own doctors or are we going to doctors sooner? Are we going to hospital sooner because we read something on the internet, we read about a symptom that we weren't aware of. So I do not fear for the future of the legal profession. I do think there'll be some reorientation, there'll be certain areas which may become where there are only gonna be a couple of people working at an extremely high level because you know a lot of the normal work can be done by an AI you know because there are areas that are pretty formulaic. But law will always be theoretical and that kind of theory based work I will say two things about. I have severe doubts whether AI can do it and I'm absolutely sure that as humans we don't want AI to do it.

**Robert Diab:**

Maybe I'll add to that only briefly to say that I agree with Jon that I don't think in a fundamental sense the practice of law will change. But I think it might change in the sense that everyday life will change, you know that the AI co-pilots or assistants will become more pervasive. It'll probably be a slow process but you know more and more of the things we do may be infiltrated by AI, you know AI will be more and more a part of a picture and that would probably carry over to law and change the way we do things but what we do, I agree with Jon that that's probably gonna be very much the same.

**Vinnie Yuen:**

Thank you so much to both of you for sharing your thoughts. Thank you so much for listening. This episode was produced by me with help from my colleague Madison Taylor. If you're a lawyer and you want more guidance on using AI tools, check out our show notes where we've linked to the Law Society's guidance on professional responsibility and generative AI. If you like this episode, be sure to us a five star rating and follow or subscribe wherever you get your podcasts.

[Podcast ends]