# AI in the Courtroom

## *A Blog Series Exploring AI and Its Intersection with the Law*

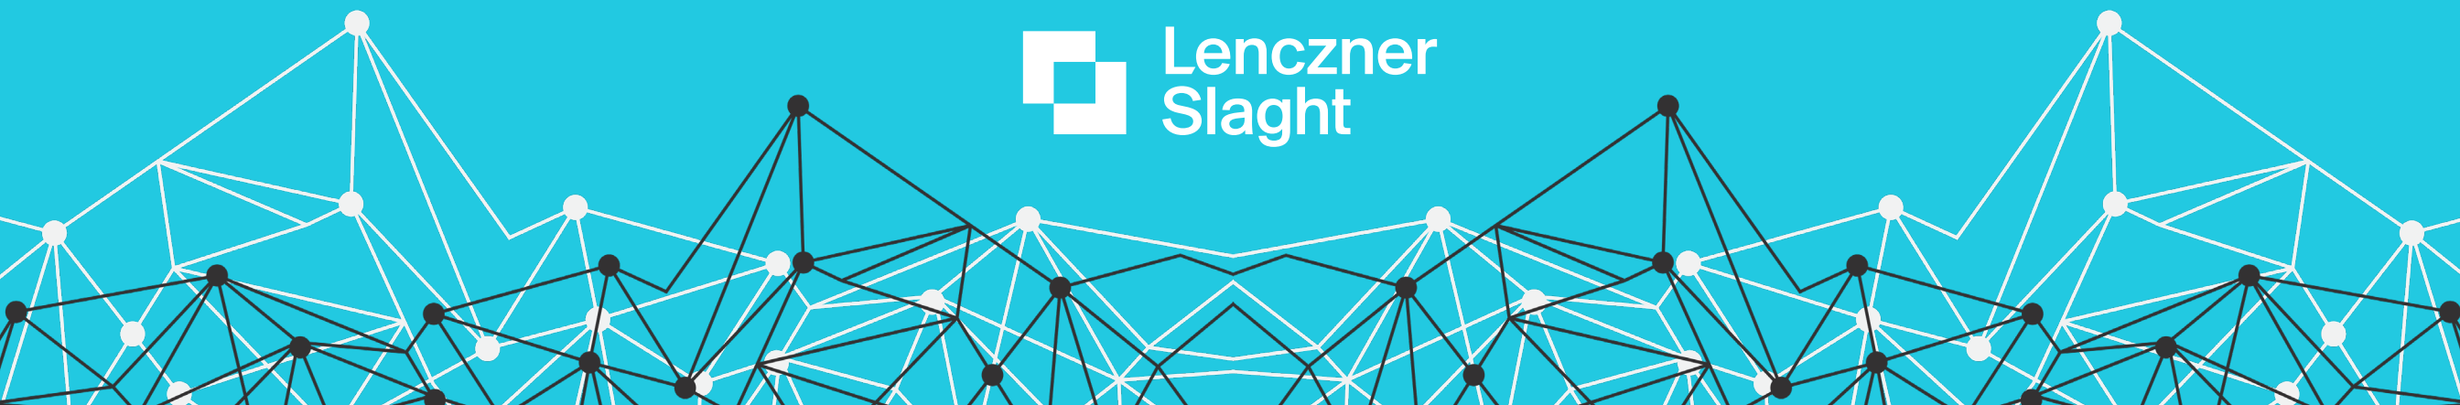

Lenczner Slaght

# Introduction

The current landscape is inundated with narratives surrounding artificial intelligence and its intersection with the law. From the New York Time's lawsuit launched against OpenAI in December 2023, to the BC lawyer reprimanded for citing fake AI-generated cases, to the lying Air Canada chatbot, the legal and mainstream media is full of stories of AI or people using AI running up against traditional legal doctrine and practice. Yet, amidst this surge of AI-related incidents, Canada finds itself grappling with more questions than answers.

Read our 5-part blog series on *AI in the Courtroom*, as we explore AI and its intersection with the law, compiled below.

# About Lenczner Slaght

Widely recognized as Canada's leading litigation practice, we have successfully represented clients' interests in some of the most complex, high-profile cases in Canadian legal history. Our lawyers are distinguished by their depth of courtroom experience, which is why people turn to us in the most difficult situations. We are bold innovators, future focused in everything we do. We've created programs that changed the way we approach advocacy, client service, and the business of law – Data-Driven Decisions, Commerciallist.com, and ReferToHer. As advocates first, we build interdisciplinary teams seamlessly, bringing together subject-matter experts to address new and complex problems, like AI. In short, we're expert litigators, prepared for the future and committed to success.

# Contents

# AI in the Courtroom: The Quest for Legal Precedents

*"The legal and mainstream media is full of stories of AI or people using AI running up against traditional legal doctrine and practice. Yet, amidst this surge of AI-related incidents, Canada finds itself grappling with more questions than answers."*

It's evident that Canadian courts will increasingly confront cases involving AI, prompting a pressing need for clarity. Here are the top ten questions (in no particular order) that we hope Canadian courts will address soon regarding AI:

- Under what circumstances does AI-generated content (e.g., deep fakes) infringe on rights of personality, privacy, and/or reputation?
- Can an AI be an author or inventor worthy of copyright or patent protection? Can it or should it be a co-author/inventor with a human? (See here for some previous discussion on this question.)
- Can generative AI companies and/or companies using their products be liable for damages caused by "hallucinations" (incorrect/fake/misleading results)?
- Is the use of a work to train an AI fair dealing? If your work is used to train an AI, what if any compensation should you receive?
- Can AIs be used in court or government decision-making, and under what conditions? (See here for a discussion of AI in administrative decision-making.)
- In areas of the law where consent is required, is disclosure that AI will be used a requirement for that consent to be informed?
- If an AI is a part of a product that causes harm to users, who is liable for that harm?
- Can an AI act as an "expert" witness?
- Is non-explainable AI a form of willful blindness?
- What standard of care should an AI/AI company be held to? What impact should voluntary codes of conduct have on such a standard?

Which questions are on your list? Do you know of a case making its way through the courts that might answer one or more of these? Please share your insights with us!

Below, Lenczner Slaght's expert litigators will explore the nexus of AI and the law in a series of blog posts. Some of them may even use the help of generative AI tools, like this blog did. We invite you to join us on this quest for legal precedents!



KEY AUTHORS AND CONTACTS

Sana Halwani
PRACTICE GROUP LEADER
416-865-3733
shalwani@litigate.com

Jordana Sanft
PARTNER
416-596-1083
jsanft@litigate.com

Jim Lepore
ASSOCIATE
416-865-2881
jlepore@litigate.com

KEY AUTHORS AND CONTACTS

Brian Kolenda
PARTNER
416-865-2897
bkolenda@litigate.com

Jim Lepore
ASSOCIATE
416-865-2881
jlepore@litigate.com

# Bars or Bytes? Exploring the Implications of a Track that Drake Might (or Might Not) Have Created

*"In this brief comment, we explore some of the circumstances in which AI-generated content might infringe on the rights of personality, identity, privacy, and/or reputation."*

Did Drake respond to an alleged feud with fellow artists, including Kendrick Lamar and Rick Ross? As reported by the LA Times, a track titled 'Push Ups (Drop and Give Me 50)' appeared online recently, taking aim at Lamar and several of his recent musical collaborators. However, this track remained unclaimed on any of Drake's official platforms, causing some to question whether this track was fan-generated using artificial intelligence. This cynicism may be justified: Drake is reportedly no stranger to having to denounce fan-generated songs, and Lamar's rumoured response was actually the work of artificial intelligence and another rapper.

Over the last few days, Drake dropped 'Push Ups' on streaming services, seemingly ending the debate regarding its source. But Drake raised a new set of interesting questions about the legal implications of AI-generated content when he later released "Taylor Made Freestyle," another Lamar diss that seemingly features AI-generated vocals from Tupac Shakur and Snoop Dogg.

**Discussion**

The above examples highlight the difficulties experienced in distinguishing authentic content from that generated by artificial intelligence. In this brief comment, we explore some of the circumstances in which AI-generated content might infringe on the rights of personality, identity, privacy, and/or reputation.

Personality

Although the tort of misappropriation of personality is "well recognized" in Canada, it is less developed than its "right to publicity" analogue in the U.S. Generally speaking, however, it arises where one's personality has been appropriated for commercial purposes (i.e., "amounting to an invasion of his right to exploit his personality by the use of his image, voice or otherwise with damage to the plaintiff"). Accordingly, so long as an individual has a valuable reputation, the use of that individual's image (in the case of an AI-generated picture or video) or voice (in the case of an AI-generated song) can be problematic. A court is likely to look at the purpose of the portrayal to determine whether it falls within the ambit of this tort (e.g., a biography where a celebrity is the subject would not expose the creator to liability, in contrast to an activity in which the celebrity is used to endorse or promote a product for commercial gain, which would). In the case of a diss track of uncertain origin that wholesale appropriates the voice of a chart-topping celebrity, the purpose of this portrayal is unlikely to provide safe harbour. Though those who followed the social media spat between Drake and Rick Ross may agree that the old adage of 'any publicity is good publicity' rings true, which raises the question whether there is any damage!

Identity / Passing Off

In addition to the misappropriation of personality, someone – like Drake – who develops content as part of their business, could also argue that AI-generated content purporting to be authentic misleads consumers. At the highest level, the tort of passing off and its codification in section 7(b) of the *Trademarks Act*, exists to protect someone from the harm arising from unfair use of their identity (e.g., pretending that a product originates from that person) and to protect the public from being misled, as to the source of particular goods or services. Much would depend on the nature of the AI-generated content in question and how it is presented; however, it is not outside the realm of possibility that such content could run afoul

Brian Kolenda
PARTNER
416-865-2897
bkolenda@litigate.com

Jim Lepore
ASSOCIATE
416-865-2881
jlepore@litigate.com



of passing off (e.g., an AI-generated song held out as from a particular recording artist, competing with that recording artist).

Privacy

Many provinces have recognized statutory or common law invasion of privacy torts. While there is some debate about the scope of such torts, Ontario courts have recognized four distinct ones which might readily apply to the misuse of AI-generated content. Intrusion upon seclusion imposes liability on a person who intentionally intrudes upon the seclusion, private affairs, or private concerns of another person, "where the invasion would be highly offensive to a reasonable person". No proof of loss or harm is required (but would be compensable, if proven). As such, to the extent AI-generated content disclosed sensitive personal details (e.g., "one's financial or health records, sexual practises and orientation, employment, diary or private correspondence") or relied on those details to generate such content, liability may arise, but "it is enough if the fact of its publication is offensive" in order to be actionable.

Unfortunately, AI-generated content is itself often used to invade an individual's privacy and can attract liability on a number of other grounds. "Public disclosure of embarrassing private facts" is actionable in Canada, such that liability may arise where artificial intelligence is used to generate, post, and amplify such content across the internet. Similarly, "publicity which places the plaintiff in a false light in the public eye" is also likely actionable in Ontario and whether on those grounds, or traditional claims of defamation, deep fake content – realistic-looking audio, video and/or images that have been altered or created using artificial intelligence – could attract liability where it is used to portray an individual in a negative light or as a tool to humiliate. Lastly, "appropriation, for the defendant's advantage, of the plaintiff's name or likeness" is also actionable in Ontario, such that malicious use of an individual's personality (as compared to commercial use described above in the context of misappropriation of personality) may also attract liability for AI-generated content in an appropriate case.

Reputation

Canadian Courts recognize several causes of action to remedy falsehoods (e.g., defamation and injurious falsehood). In the commercial context, section 7(c) of the *Trademarks Act* prohibits certain false or misleading statements against competitors. Where registered trademarks are involved, section 22 of the *Trademarks Act* prohibits certain uses of well-known marks (or indicia linked thereto) in a manner that depreciates its goodwill. As such, many of the examples canvassed above for Personality, Passing Off and Privacy – which by their nature constitute a falsehood – may also attract liability under reputation-related torts.

*"There is a need for those in the creative and tech industries to understand the legal implications of AI-generated content."*

**Takeaways**

There is a need for those in the creative and tech industries to understand the legal implications of AI-generated content. The questions of AI-generated content raised in the ongoing rap feud between Drake and Lamar highlight broader challenges likely to come before our Courts – contending with AI-generated content that engages several aspects of the law at any given time, from personality to privacy. If AI sets the rhythm for tomorrow's tracks, the law must keep pace – without skipping a beat.

**Update**

OpenAI recently introduced a voice for its ChatGPT product that some people say sounds "eerily similar" to Scarlett Johansson's voice. This follows Johansson's refusal of an offer from OpenAI to use her voice. We continue our discussion on AI-generated voice and explore the legal implications of this situation here.

# On the Horizon: Legal Complexities Intersecting Generative AI, Class Actions, and IP Law

*"As generative AI becomes more mainstream, people will look to class action proceedings to address their grievances."*

The multifaceted nature of generative AI is bound to create legal complexities at the intersection of intellectual property law and class actions, as this emerging technology disrupts not only the tech landscape but the legal one too.

Fostering innovation is a core tenet of intellectual property law in Canada. However, policy interests seek to balance the furtherance of technology and creativity with protecting the public. Class actions offer a process for advancing public interest by allowing representative plaintiffs to advance claims on behalf of an entire class of people. Class actions advancing consumer rights are commonplace. As generative AI becomes more mainstream, people will look to class action proceedings to address their grievances. As this happens, patent, copyright, trademark or other IP related allegations are also likely to become more present in such actions.

Consider, for example, an AI generated consumer product that fails to perform as intended or causes harm due at least in part to a patented technology. Affected individuals may seek recourse through a product liability class action. Similarly, disputes over data scraping or use of copyrighted materials to train AI or generate works are on the horizon. Creative works such as art and literature are already being used to train AI systems to create new works. Indeed, issues like these are already starting to appear in legal battles in the United States (see for example: *PM v OpenAI LP*; *JL v Alphabet Inc*; *Andersen v Stability AI Ltd*; and *Getty Images (US), Inc v Stability AI, Inc*.

**Discussion**

Liability

A significant issue that class actions pertaining to generative AI will raise is establishing liability and accountability for the harm caused by AI. This involves grappling with questions about whether liability should and could fall on an AI system, its developers, users, or owners, and how to allocate responsibility between them. Liability becomes particularly challenging when multiple parties are involved in the development and deployment of the generative AI technology.

Proving harm and causation brings unique challenges in the context of generative AI. Demonstrating how the content created by generative AI harmed class members will be among the hurdles that need to be overcome if liability is to be made out. Additionally, establishing the causal link between an AI system and the alleged harm will pose its own challenges, likely requiring a host of experts with legal and technology appropriate backgrounds to advance strategies and present a cohesive case.

Disputes over IP ownership and inventorship could also arise in the context of a generative AI class action. The patent holder or the class could dispute liability on the basis of whether the generative AI is liable and to what extent, if any, its owners, creators, and authors/inventors may be held accountable. These legal battles could also assert complex issues of IP rights and enforcement like patent validity and infringement into the class action arena, with the further potential of impacting remedies like quantum of damages. Addressing such issues would require a deep understanding not only of generative AI and class action strategy but also of IP law, pushing the legal arena into unchartered territory.

Standing

Standing is another obstacle to overcome in class action cases directed at generative AI issues. For example, it remains unclear whether unauthorized use of the copyright works or personal data in AI models results in a legally recognizable harm, and if it is does, whether this constitutes an injury that is sufficient for a plaintiff to pursue their legal recourse and theory



KEY AUTHORS AND CONTACTS

Jordana Sanft
PARTNER
416-596-1083
jsanft@litigate.com

Paul-Erik Veel
PRACTICE GROUP LEADER
416-865-2842
pveel@litigate.com

of the case. Would an exception to IP infringement such as fair use dispel an allegation of wrongdoing or harm? Depending on the circumstances, the use of copyrighted materials only for training purposes may not involve impermissible copying or substantial reproduction for commercial purposes as traditionally contemplated under copyright law. The viability of such claims remains to be seen.

*"While liability risk should not dissuade companies from working in the generative AI space, companies should be mindful of the liability risks posed and take steps to minimize them."*

Reverse Class Actions

Another interesting twist may be the rise of reverse class actions. In a conventional class action, a representative plaintiff sues a defendant on behalf of a class. In a reverse class action, an individual plaintiff sues a group of defendants who are alleged to have engaged in the same wrongful conduct. These reverse class actions require a representative defendant. There are several recent examples in the area of IP albeit, not involving AI (see *Voltage Pictures v Salna* and *Seismotech IP Holdings Inc v John Does*).

These cases may unwittingly have laid the foundation for increase in class action proceedings before the Federal Court where generative AI and IP law are in issue. Although such actions could be brought in superior courts, the Federal Court with its strong IP capability and the ability to invalidate a patent in rem, creates an interesting choice of venue depending on the specifics of the issues in dispute.

The intersection of generative AI, class actions, and IP law presents a host of intricate issues that require careful consideration and expertise. To this end, there is a need for collaboration between experts and legal specialists working towards a holistic strategy that promotes innovation, advances IP rights, and grapples with class action standing and liability, among other issues, relating to the generative AI ecosystem.

**Practical Tips**

There remains a myriad of unresolved legal issues in this space, and it will be challenging for players operating in this area to take steps to completely avoid any risk of liability. Novel and disruptive technologies virtually always present some risk of liability by their nature. That is a risk that many players choose to accept in order to build novel technologies and products.

That being said, while liability risk should not dissuade companies from working in the generative AI space, so too should companies be mindful of the liability risks posed and take steps to minimize them. For example:

Monitor the Landscape and Identify Best Practices

Participating in industry forums and knowledge sharing in the evolving landscape can provide important insight, help set the stage of new industry norms, and provide a means of mitigating risk. By proactively engaging in discussions assessing and addressing social, ethical, and legal implications, additional innovation fostering safeguards may be identified and adopted.

Further, by identifying and tying to address possible risks, one can minimize the potential for liability and be better prepared to navigate class action and IP issues.

Engage Experts

When a potential issue arises, move quickly to get advice as to how to handle it. Engage with legal and technology experts early in the pre litigation and litigation context to obtain guidance on the complex legal and factual issues in your industry facing innovation and commercialization.



KEY AUTHORS AND CONTACTS

Jordana Sanft
PARTNER
416-596-1083
jsanft@litigate.com

Paul-Erik Veel
PRACTICE GROUP LEADER
416-865-2842
pveel@litigate.com

KEY AUTHORS AND CONTACTS

Sana Halwani
PRACTICE GROUP LEADER
416-865-3733
shalwani@litigate.com

Jonathan Chen
PARTNER
416-865-3553
jchen@litigate.com



# AI Competence in the Courtroom: Four Things Judges Need to Understand Now About AI

As artificial intelligence continues to permeate every aspect of our lives, legal challenges involving AI will proliferate. Parts 1 to 3 in our series explored many of these potential questions. AI will create new legal problems and change the texture of old ones. As always, the judiciary, with the assistance of counsel, will assume a pivotal role in navigating this landscape.

Grappling with technology is nothing new for judges, but the combination of complexity, rapid evolution, and expected ubiquity of AI means that judges are at risk of getting it very wrong, very easily. With that in mind, we provide basic answers to four questions judges need to understand about AI before grappling with any case involving this technology.

**1. How is AI different from other sophisticated software?**

AI differs from other sophisticated software in several fundamental ways. AI is capable of learning, adapting, and performing complex tasks autonomously, distinguishing it from traditional software.

- **Learning and Adaptability:** AI systems, particularly those using machine learning, can learn from data and improve their performance over time without being explicitly programmed for each task. For instance, a machine learning model can improve its accuracy in predicting outcomes as it is exposed to more data. AI systems can adapt to new situations by retraining on new data. This adaptability allows AI to function in dynamic environments and solve complex, variable problems.

  On the other hand, traditional software follows a predefined set of instructions written by programmers and is less or not adaptable at all. It performs tasks exactly as programmed and does not improve or adapt unless explicitly updated by developers. Changes in its functionality or environment typically require manual code updates by developers.

- **Data-Driven:** AI relies heavily on data for training and decision-making. The performance of AI models often correlates with the quality and quantity of data they are trained on. One such quality issue is data bias which we discuss below. Traditional software is not inherently data-driven. While it can process data, its functionality is more dependent on the specific code and logic defined by programmers rather than on data analysis and learning.

- **Decision-Making and Autonomy:** AI can make decisions based on data analysis and pattern recognition. It can handle unstructured data (like images, text, and voice) and make decisions that mimic human reasoning. AI systems can operate with a high degree of autonomy, performing complex tasks with minimal human intervention. Traditional software makes decisions based on fixed logic and predefined rules; it lacks the flexibility to interpret unstructured data. It requires ongoing human input and supervision, executing tasks based on specific user commands.

- **Human-Like Interaction:** AI enables more natural interactions with humans through technologies like chatbots, virtual assistants, and voice recognition systems. These systems can understand and generate human language to some extent. Traditional software interactions are typically more rigid and limited to predefined interfaces and commands, lacking the nuanced understanding of human language.

As addressed below, these differences create some of the thorny problems that judges will have to grapple with as they address cases involving AI.

**2. What is AI bias, and why does it exist?**

Just like humans, AI can be biased too. It is the human that develops and trains the artificial intelligence model. AI bias refers to artificial intelligence models that produce results which reflect human biases. These biases can in turn perpetuate historical social inequities. Take, for example, an AI recruiting tool that unintentionally favors candidates with a certain background or interest or of a specific gender. The result is a discriminatory one (even if unintended).

Bias can seep into artificial intelligence in several ways. Two common examples are through the training data and the algorithm. An AI system is only as good

Sana Halwani

PRACTICE GROUP LEADER
416-865-3733
shalwani@litigate.com

Jonathan Chen

PARTNER
416-865-3553
jchen@litigate.com



as the data input. Artificial intelligence models learn to make decisions based on that data. For example, generative AI models are built to generate text (which word should come next?) by relying on probabilities based on the dataset the model was trained on (which word usually comes next in this context?). If the dataset itself is incomplete (e.g. a certain variable is over or underrepresented in the dataset), skewed or outdated, then the probabilities and therefore the predictions will reflect those limitations. The algorithm employed can also be tainted by the developer who may inject personal preferences or weight certain attributes more heavily than others.

So, even with the best of intentions, the artificial intelligence model may produce a biased result that is perpetuated and amplified by someone also with the best of intentions. Awareness of this reality is imperative before evaluating any allegations relating to AI models, accepting results from AI models (e.g., through expert evidence), or incorporating artificial intelligence into legal decision making.

**3. What is the difference between explainable and non-explainable AI, and why does it matter?**

Explainability is about building trust in the artificial intelligence model. When we are using AI models to make predictions, there is a natural tendency for lawyers to ask how it came to that result. Developing tools and processes to understand that result is explainability.

Explainability is tied to the concept of responsible use of artificial intelligence. Whether used in a business or for legal matters, the artificial intelligence model should not be a "black box". ChatGPT, for example, is a black box because you do not know how it came to the conclusion that it did. As end users, we need to know that the model is competent, trustworthy, safe to use, up to date and accountable. To be accountable, the model must be understandable and able to be subject to human oversight and scrutinization. This in turn allows the user (or the judge) to determine whether the model meets, for example, background requirements (e.g., company policies, regulatory standards, or practice directions) or has been tested and validated. Explainability therefore imbues the result with reliability. Judges will have to grapple with what level of reliability is required in a given situation.

In addition, if a situation where harm is alleged to have been caused by a "faulty" AI, whether that AI is explainable may affect the ability of the Court or parties to evaluate fault. Self-driving cars are often used as the example here. If a self-driving car gets into an accident and its decision-making is impugned, how will a Court evaluate fault if the decision-making cannot be explained?

**4. Can generative AI lie, and can a human tell if this is happening?**

Generative AI models like ChatGPT can produce outputs or answers that are incorrect or misleading. While an AI may not have the intent for such misleading content to be called "lying", the impact may be similar.

Courts have already started to grapple with such fabrications, sometimes called "hallucinations." For example, earlier this year, the Supreme Court of British Columbia issued a decision addressing a notice of application containing fabricated legal authorities that had been "hallucinated" by ChatGPT. The lawyer who included them gave evidence at the hearing that she did not know ChatGPT could generate fake authorities.

Generative AI models create text based on the statistical likelihood of word sequences. This means they can produce plausible-sounding but incorrect or nonsensical responses if the data suggests such patterns. Generative AI models lack the contextual understanding that a human might have. If a user inputs a prompt that is ambiguous or open to interpretation, the AI might generate a response that might fit the prompt but may not be factually accurate. The AI does not understand the context in the way humans do.

Determining whether an AI is providing accurate information or producing hallucinations can be extremely challenging for humans. This difficulty is compounded if the AI is not explainable (which is the case currently for most if not all iterations of generative AI models). Understanding the limitations of an AI (including the data on which it was trained); fact-checking; and consulting multiple sources not just AI sources, are three strategies that can help. But to employ any strategies, a Court will need to know if generative AI was used, whether in the context of a lawyer's brief or the facts of a case.

This instalment of our *AI in the Courtroom* series strives to provide judges with a basic understanding of some technical background and issues that may arise when AI is used in the courtroom. Judges do not need to become data scientists or coders to manage the use of AI in the courtroom, or to evaluate cases involving AI. It is important, however, for judges to at least be aware of how the AI before them was developed, how it works, its application to the particular case, and the risks and implications. When these key issues are kept in mind, judges will be able to play their role as gatekeeper and properly assess when to ask questions, what questions to ask, and at what level of detail. Counsel would of course be wise to prepare responses in advance to effectively assist the Court.

# AI Here, AI There, AI Everywhere: Practical Challenges Litigating in an AI World

In the final instalment of our *AI in the Courtroom* series, we explore practical challenges that may arise when litigating in an AI world, and within the current framework of the *Rules of Civil Procedure*, Practice Directions, and common law. While the law is not entirely unequipped to deal with these challenges, evolution in the *Rules* and common law will likely be necessary as AI becomes more commonly used by various participants in the litigation process.

### 1. Risk of AI-generated Evidence

Civil litigators habitually challenge the admissibility of evidence on the basis of its relevance or reliability. Authenticity—whether a document or other piece of evidence actually is what it purports to be—is less frequently the basis of an objection. It is uncommon in civil litigation to dispute that an email was sent by the person in the "from" line, received by the person in the "to" line, at the time indicated on the message. But where there is concern that evidence was generated by AI and is not "real", the way to challenge the admissibility of that evidence is through objecting to its authenticity.

The problem of deep fakes will arise in litigation when there is a dispute between the parties about whether a specific piece of evidence, such as a text message, voice note, or video, is real or the product of AI. A dispute about deep fakes is at its core a dispute about authenticity. While expert evidence might be able to resolve this question, not every case is going to involve experts. And unlike forgery, which previously required some level of skill to do well, anyone with a computer and access to the internet can now create a deep fake if they choose. How, then, can Courts address this problem?

The threshold for authentication of evidence is low. In *R v CB*, the Court of Appeal for Ontario held that to the extent there is a dispute about whether the evidence has been tampered with, there must be an "air of reality" to the claim about tampering, and in any event the issue of tampering would likely go to the weight of the evidence rather than its admissibility given the low threshold for authentication.

These principles may have worked well in an era before generative AI made it easy to fake a text message conversation or even a voice recording. However, with the advent of generative AI, there is a risk that evidence will be admitted even where there is a serious dispute about whether it was tampered with or created by AI because the threshold for authenticity is so low. Significant trial time may then be wasted adducing evidence about the alleged tampering and/or "deep fake" nature of the evidence, only for that evidence to go to weight rather than keeping the tampered with or deep fake document out of the court record in the first place.

The risks and problems posed by deep fakes in the era of generative AI is real. But wariness of deep fakes has another, equally challenging problem for litigators: what happens when a party knows a document is real, but alleges it is a deep fake in an effort to discredit that evidence or the other party? The only remedy to this problem currently available to Ontario courts is a heightened costs award. In *Jurrius v Rassuli*, a family law dispute, the father alleged that a photograph of a replica gun strapped to the child's crib included in the applicant mother's materials was "doctored" or "photoshopped". On cross-examination at trial, he admitted he had in fact strapped the replica gun to the child's crib and knew the photograph in the mother's materials was valid. The father's misrepresentation about the photograph was criticized in strong terms and an important basis for the court's award of full costs. But a costs award made after the litigation is over is small comfort, given the seriousness of the allegation that evidence is fake (whether a deep fake or otherwise).

### 2. Expert Evidence Dependent on AI

Experts play a critical role in complex cases before the courts, but they can only play that role well if they are properly qualified and abide by their duties to the court.

As AI tools proliferate, courts will have to grapple with whether expert opinions that rely on AI or were generated by AI should be admitted as evidence. At the very least, the usual rules of evidence would apply: the four criteria for the admissibility of expert evidence are:

(1) relevance;

(2) necessity in assisting the trier of fact;

(3) the absence of any exclusionary rule; and

(4) proper qualification (*R v Mohan*).



KEY AUTHORS AND CONTACTS

Sana Halwani
PRACTICE GROUP LEADER
416-865-3733
shalwani@litigate.com

Meghan Bridges
PARTNER
416-749-3974
mbridges@litigate.com

These criteria provide significant discretion to the Court to, for example, exclude expert evidence on the basis that a generative AI model came to the "opinion" reported by the expert. Such "opinions" would – arguably – not have come from the qualified expert and would not be admissible. In contrast, an expert using an AI tool to assist in their analysis, would pose less concerns.

While Canadian courts have started to publish practice directions that address the use of AI by counsel and the Court, none have – to the authors' knowledge – yet addressed the use of AI by expert witnesses. For example, the Federal Court's Notice states: "This Notice requires counsel, parties, and interveners in legal proceedings at the Federal Court to make a Declaration for AI-generated content (the "Declaration"), and to consider certain principles (the "Principles") when using AI to prepare documentation filed with the Court." There is no mention of experts.

It would therefore appear that generative AI and other AI tools can be used by experts to generate expert reports and inform their opinions without disclosure being required. Arguably, existing rules and codes of conduct may apply to prevent such situations in certain circumstances, for example, by requiring an expert to disclose the methodology used for any testing he or she conducted. But these kinds of requirements do not explicitly apply to AI and are open to interpretation.

Given the centrality of expert opinions to certain kinds of cases, addressing the use of AI by experts will be critical to ensuring the fairness and transparency of the litigation process.

### 3. Use of AI by Decision-Makers

Judges and administrative decision-makers will certainly not be immune from the lure of using AI in generating decisions. And neither should they be, so long as safeguards are in place to protect from bias and ensure procedural fairness. Court systems in Ontario and across Canada are in crisis, and AI may be part of a solution to that crisis. This is nothing new. The legal profession has (slowly, begrudgingly) embraced technology in the last few decades – from word processing, to legal research databases, to e-discovery tools – resulting in great gains of efficiency.

Thus far, courts are taking it slowly with AI. For example, the Federal Court has addressed this issue in its "Interim Principles and Guidelines on the Court's Use of Artificial Intelligence" stating that it:

> *"will not use AI, and more specifically automated decision-making tools, in making its judgments and orders, without first engaging in public consultation."*

This is a reasonable starting stance. The public needs to be confident that its judicial and quasi-judicial decision-makers are not delegating their responsibilities away. One of our colleagues has explored the impact of AI on administrative law and procedural rights more fully here.

Decisions from administrative decision-makers have already started to be challenged on the basis that the decision-maker used an AI tool. For example, in Haghshenas v Canada (Citizenship and Immigration), the applicant argued that a decision made by an immigration officer with respect to a work permit was unreasonable and not procedurally fair as it was reached with the help of an AI system called Chinook.

We pause here to say that we question whether the Court should have accepted that Chinook was properly characterised as an AI tool. In fact, Immigration, Refugees and Citizenship Canada's statement on "Chinook Development and Implementation in Decision-Making" states that:

> *"Chinook is a tool designed to simplify the visual representation of a client's information. It does not utilize artificial intelligence (AI), nor advanced analytics for decision-making, and there are no built-in decision-making algorithms."*

Regardless, the Court proceeded as if an AI tool had in fact been used in the decision-making process.

In dismissing the application, the Court determined that the decision was made by the officer, not by Chinook, though the officer did consider input compiled by the AI. The Court highlighted that the use of AI was irrelevant to the judicial review application because the officer ultimately made the administrative decision. The Court concluded on this issue with:

> *"Whether a decision is reasonable or unreasonable will determine if it is upheld or set aside, whether or not artificial intelligence was used. To hold otherwise would elevate process over substance."*

While this is an attractive framing, it fails to acknowledge that reasonableness review may be hampered by the use of AI tools, for example, if their results are not explainable (see our previous blog which describes explainable vs non-explainable AI, and why judges need to understand the difference). As Courts' understating of AI becomes more nuanced, we expect to see more detailed and nuanced guidance on when use of AI in decision-making is acceptable and when it is not.

### Takeaways

Whether addressing the possibility of deep fake evidence, AI-generated expert opinions, or robot decision-makers, what the cases described above tell us is that counsel and the Courts must remain vigilant in ensuring that no part of the litigation ecosystem is abdicating their responsibilities to AI, even if AI is here, there, and everywhere.



## KEY AUTHORS AND CONTACTS

**Sana Halwani**
PRACTICE GROUP LEADER
416-865-3733
shalwani@litigate.com

**Meghan Bridges**
PARTNER
416-749-3974
mbridges@litigate.com

@LencznerSlaght
/LencznerSlaght