Artificial Intelligence in the Practice of Law (forthcoming 2024)

The Sedona Conference Working Group 7
(Sedona Canada)

Draft for review and comment. Further distribution restricted.

ARTIFICIAL INTELLIGENCE IN THE PRACTICE OF LAW

Contents

I. INTRODUCTION

"Artificial Intelligence" ("AI"), a term that once existed primarily in the realm of science fiction, has swiftly transitioned into a pervasive reality in many sectors, including the legal domain. The implementation of AI is reshaping the traditional contours of the legal profession, offering a myriad of opportunities while simultaneously presenting novel challenges. This primer aims to provide a broad understanding of the role AI plays in the practice of law and its potential impact moving forward.

"AI" refers to a collection of technologies that replicate human intelligence by performing cognitive tasks, such as perceiving, learning, reasoning, problem-solving, and understanding and generating language. Its application in law, while still in the early stages, has already begun to influence various aspects of legal practice, from document review and legal research to contract analysis and prediction of legal outcomes.

By way of example, in the realm of document review, AI can quickly sift through vast quantities of information, identifying and highlighting key pieces of evidence, saving considerable amounts of time and money. Legal research, a fundamental aspect of legal practice, has also experienced a significant transformation with AI. Advanced legal research platforms are now capable of providing more targeted and relevant case law, statutes, and secondary sources in a matter of seconds, in response to a natural, plain-language query, a task that in the past might have taken hours, if not days of search. AI's use in contract analysis has revolutionized the way lawyers review and draft contracts. AI tools can identify standard clauses, flag missing clauses, and even suggest language based on the lawyer's prior work product, thereby increasing efficiency and minimizing human error. Finally, AI can analyze past judgments and rulings to forecast the potential outcome of a case, thereby assisting lawyers in formulating their legal strategies.

While these advancements offer immense potential, they also raise a host of ethical and practical challenges. As AI begins to automate aspects of the legal process, questions arise about the future role of lawyers. There are concerns about data privacy and confidentiality, given the large volumes of sensitive information AI systems handle. In addition, the use of AI in predicting legal outcomes and other issues raises questions about bias and fairness, as these predictions are based on past decisions or data, which may have been influenced by inherent biases.

This primer delves into these and other issues, providing an overview of how AI is transforming the practice of law. It aims to explore the various types of AI tools currently being used, their implications for lawyers, law firms, clients, and judges, as well as potential future developments in this field. This primer will also discuss ethical considerations that arise from the use of AI in law and how the legal profession can navigate these challenges.

In essence, the intersection of AI and law presents a profound paradigm shift in the way legal services are and will be delivered. Whether one views this as a threat to the traditional legal profession, or an opportunity for growth and innovation, the fact remains that AI's impact on the

practice of law will be massive and far-reaching. This primer, therefore, seeks to shed light on this evolving landscape, offering insights that will be invaluable for legal practitioners, clients, judges, law students, and anyone interested in the future of legal services.

A. Scope

The focus of this paper is on "Narrow AI," also known as "Weak AI." "Narrow AI" is the branch of AI systems that are designed to perform a specific task and operate under a limited set of constraints. Unlike General or Strong AI, Narrow AI does not possess the ability to understand, learn, or apply knowledge beyond the specific task for which it has been programmed.

Examples of Narrow AI are prevalent in our everyday lives. These include Siri and Alexa, which use voice recognition to respond to user commands, recommendation systems on e-commerce sites like Amazon, which suggest products based on a user's browsing history, and AI in video games that adapt to a player's behavior.

While Narrow AI can mimic human intelligence for specific tasks, it does not possess true understanding or consciousness. It operates based on pre-programmed instructions or learned patterns and does not have the ability to tackle tasks outside its designed scope. Despite its limitations, Narrow AI has made significant progress in numerous fields, including healthcare, finance, transportation, and entertainment.

This paper will not provide a lengthy discussion of "General AI," also known as "Strong AI," which remains a theoretical concept and the subject of considerable research and debate. "General AI" refers to the branch of AI that would be able to understand, learn, adapt, and implement knowledge on a wide range of tasks, in a manner akin to a human being. Unlike Narrow AI, which is designed to perform a specific task, such as voice recognition, General AI would be able to perform any intellectual task that a human being could do.

Similarly, this primer will not cover "Superintelligence," which refers to an AI intellect that is significantly more advanced and capable than the human mind. It encompasses superior problem-solving skills, creativity, and knowledge in every field, such as scientific reasoning, social skills, and general wisdom. At present, superintelligence resides solely in the realm of science fiction.

General AI and Superintelligence are briefly addressed in Section VII.

B. Automation vs. Augmentation

Automation and augmentation, though closely related concepts, have distinct implications, particularly in the context of AI and its applications.

"Automation" refers to the process of using machines, robotics, or AI to perform tasks that were previously performed by humans, but can now be performed mechanically, with minimal or no human intervention. It aims to increase efficiency, speed, and accuracy, and is often used for

repetitive, routine tasks. Examples of automation include assembly lines in manufacturing plants, self-checkout systems in supermarkets, or email spam filters.

On the other hand, "augmentation" refers to the use of technology to enhance or improve human capabilities, rather than replace them. Augmentation tools are designed to assist humans in performing tasks more effectively and efficiently, by providing insights, decision-making support, or enhanced abilities. Examples of augmentation include AI-driven document review analysis tools that help lawyers make better decisions about collections of documents or GPS systems that enhance our natural navigation abilities.

In essence, while automation seeks to *replace* human effort, augmentation aims to *enhance* it. Both have significant roles to play in the future of legal services and the legal profession, and the balance between them is a key consideration in the design and implementation of AI systems.

II. DIFFERENT TYPES OF AI

A. Expert Systems

An "Expert System" is a computer system that emulates the decision-making ability of a human expert using AI techniques. Expert systems are designed to solve complex problems by applying a deductive logical structure consisting of dynamically created if–then rules (i.e., a list of conditions and responses) rather than through conventional procedural (i.e., step-by-step) code. Expert systems were among the first truly successful applications of AI, with the first systems appearing in the late 1970s.

An expert system is divided into two subsystems: the "inference engine" and the "knowledge base." The "inference subsystem" applies rules to the known facts to deduce new facts. Inference engines can also include explanation and debugging abilities.

Originally, the "knowledge base" subsystem was composed of pre-defined facts and rules. With the advent of machine learning, big data, and data-mining techniques, modern expert systems can now incorporate new knowledge more easily and thus readily update themselves. Such systems can generalize from existing knowledge to better deal with vast amounts of complex data.

B. Machine Learning

"Machine Learning," or "Computer-Assisted Learning," is used for solving problems where the cost to develop algorithms by human programmers would be prohibitively expensive or time consuming. Machine-learning systems solve problems by having the computer system build a model of the problem without needing to be explicitly instructed what to do, step by step, by any human-developed algorithms. In other words, machine-learning systems infer the rules necessary to accomplish a task without having been specifically programmed or taught those rules.

Machine-learning approaches have been applied to large language models ("LLMs"), computer vision, speech recognition, email filtering, legal document review, agriculture production, as well as medical research and diagnosis, among other areas.

Machine learning approaches are traditionally divided into four broad categories, as follows:

1. Unsupervised Learning;
2. Supervised Learning;
3. Reinforcement Learning; and
4. Deep Learning.

"Unsupervised Learning" takes a set of data and looks for some type of structure, like groupings or clustering of data points. Instead of responding to human feedback, as is the case with supervised and reinforcement learning, unsupervised learning automatically identifies commonalities and/or anomalies in the data to build its model. An example of unsupervised learning is document

clustering capability in eDiscovery tools, which group together documents with similar topics or traits.

"Supervised Learning," also referred to as "Active Learning," builds a mathematical model from a set of labeled training data (i.e., positive and negative exemplars), through an iterative approach, whereby humans label or categorize data and test the system's ability to learn to categorize or label uncategorized or unlabeled data. For example, the training data set could be a set of documents to be reviewed by a producing party in a litigation. The supervised learning system would present a small subset of the documents to human lawyers and ask them to indicate if the contents of the document are responsive or unresponsive to the requesting party's Requests for Production. The system would then use the contents of the documents and the human coding decisions to build a decision model. It would then continue to present subsets of the documents to the lawyer for coding, updating and refining the model with each iteration. When the model's decisions and the lawyer's decisions match, the model is said to have "stabilized" and can be relied upon to categorize the remaining documents. This process has typically been referred to as "TAR 1.0."

"Reinforcement Learning" is an area of machine learning that involves teaching a system how to behave or take action using a reward and punishment structure. In essence, reinforcement learning rewards the system when it makes a correct decision, and penalizes the system when it makes an incorrect decision. Examples of reinforcement-learning algorithms are those used in autonomous vehicles, or in systems used to play games against human opponents.

"Deep Learning" involves a series of layered algorithms, starting with an input layer, followed by some number of hidden layers, and then an output layer. Each layer performs a different function and passes certain information on to the next layer. Often referred to as an "artificial neural network" because the logical structure and functioning is presumed to be similar to a biological brain, with interconnected nodes and synapses, deep learning enables machines to perform complex tasks and make accurate decisions without help from humans. Deep-learning system are far more capable than standard machine-learning models. Like supervised and reinforcement learning, deep learning requires training to ensure accurate results. Examples of deep learning are Computer Vision (used for interpreting and labeling images) and Generative AI (used for developing new creations based on existing data).

C. Natural Language Processing

Natural Language Processing ("NLP") is the branch of AI focussed on providing computers with the ability to interpret, understand, and generate language. Since virtually all legal practice involves human language, NLP plays a significant role in many legal applications. This can range from interpreting text from pre-existing data, such as case law, statutes, briefs, or evidence, to generating text from deposition recordings. NLP allows applications to accurately and efficiently

analyze massive volumes of legal text so they can be searched, summarized, translated, or analyzed by legal practioners.

At a very high level, NLP involves the pre-processing of data, training a model, and using the model for a specific purpose (e.g., as part of a legal application). Pre-processing of data typically includes breaking sentences into smaller units of words or phrases, identifying the root form of words (e.g., "teaching" into "teach"), and removing stop-words that do not add specific meaning (e.g., "a" or "the"). A machine-learning algorithm is then applied to the pre-processed data to train an NLP model. The model can then be run over the data within a live or production environment.

D. Computer Graphics

Computer graphics is a field of computer science in which computers apply various algorithmic and mathematical processes to create or manipulate images. Computer graphics is applied in many areas, including video games, animated films, visual effects, computer aided design ("CAD" drawings), computer aided manufacturing ("CAM"), simulations, medical imaging, and information or data visualization.[1]

Recent advances in AI now allow computers to generate images based on the text submitted by a user (e.g., "a cute white kitten dreaming of goldfish"). These text-to-image applications are a form of Generative Artificial Intelligence ("Gen AI"), whereby a computer can generate new content, such as text, audio, images, or video. A "Diffusion Model" is a form of Gen AI that can generate new images that are similar to those on which the model was trained. By adding random noise to the trained images, the model learns how to remove the noise and to construct the image described by the user.[2]

E. Computer Vision

Computer vision is a related area of computer science that deals with images and videos. Unlike computer graphics, where computers create or manipulate images and videos, computer vision uses computers to identify objects and people within those media. Generally, a computer-vision application analyzes an image or video in three steps. First, a sensor device, such as a camera or medical-imaging device, captures the image. Next, the image is passed to an interpreting device that breaks down the image into patterns and compares those patterns against the library of patterns on which the application has been trained. Finally, a user submits a request about the image (e.g.,

[1] STEVE MARSCHNER AND PETER SHIRLEY, FUNDAMENTALS OF COMPUTER GRAPHICS 1-3 (5th ed 2015).

[2] Victor Dey, Venture Beat, October 26, 2022, *How diffusion models unlock new possibilities for generative creativity*, https://venturebeat.com/ai/how-diffusion-models-unlock-new-possibilities-for-generative-creativity.

"Is the person in the image a company employee?") and the interpreting device provides an answer based on its pattern analysis.[3]

Computer vision and AI are applied in developing autonomous vehicles. For example, Tesla vehicles come with an array of sensors that collect data from the surrounding world. The onboard computer uses deep learning to train on the sensor data. This training enables the computer to recognize different objects within the world (e.g., a fire hydrant) by comparing the collected image against millions of example images collected from other Tesla vehicles or the Internet (i.e., other images of fire hydrants).[4]

In addition to identifying objects, computer vision and AI are also used for searching and classifying images. For example, Amazon's Rekognition is an image-recognition service that detects objects, scenes, activities, landmarks, faces, dominant colors, and image quality. When searching for faces, Rekognition applies deep learning to identify faces within an image. When a face is identified, the application applies a rectangular frame around the face along with a confidence score indicating the likelihood of the match.[5]

F. Speech Recognition

Speech-recognition programs process human speech into written format. These programs are often referred to as "automatic speech recognition" ("ASR"), "computer speech recognition," or "speech-to-text." Speech recognition is used across many industries, including automotive, technology, legal, healthcare, sales, and security. Typically, these applications consist of several components, such as the speech input (e.g., an audio recording or a user's voice), feature extraction, feature vectors, a decoder, and text output.[6]

For example, transcription services are commonly used in the legal industry and allow lawyers to take a recording (e.g., of a deposition or client interview) and upload the recording to a transcript-service site. The lawyer can then download a computer-generated transcript that is at least 90% accurate.[7] Speech-recognition programs also use AI, allowing the program to learn and improve transcription accuracy as more speech is processed. These programs may use neural networks to analyze and train on various aspects of language, such as grammar, syntax, and/or structure.[8]

---

[3] Microsoft Azure Resources, *What is Computer Vision?*, https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-is-computer-vision/.
[4] Ben Dickson, TechTalks, September 17, 2018, *The challenges of teaching driverless cars to see the world*, https://bdtechtalks.com/2018/09/17/self-driving-cars-ai-computer-vision.
[5] Amazon Rekognition Image (https://aws.amazon.com/rekognition/image-features).
[6] IBM, *What is speech recognition?*, https://www.ibm.com/topics/speech-recognition.
[7] REV, *How to Automatically Transcribe Your Audio and Video to Text*, https://www.rev.com/blog/resources/how-to-automatically-transcribe-audio.
[8] IBM, *What is speech recognition?*, https://www.ibm.com/topics/speech-recognition.

G. Audio Search

In eDiscovery, speech-recognition programs are also used to search voicemails or call-center recordings. Once these recordings are converted to a transcript, search and analytics tools can be used to identify relevant portions of the recording. For example, Intelligent Voice is a tool that is available within Relativity's eDiscovery application. This tool allows a user to interact with an audio player that plays the audio file alongside the text transcription. The user can search the transcript, review a summary of the main topics discussed, and jump to various parts of the conversation.[9]

H. Audio Generation

In addition to converting human speech into text, computer applications can convert text to audio, such as human speech, music, and ambient environmental sounds. Text-to-speech ("TTS") or speech-synthesis technology is used in virtual assistants, whereby a computer generates the virtual assistant's synthetic voice based on written text. These virtual assistants humanize the transaction and can assist visually impaired users, as well as those that are vocally challenged.

To generate a human-like voice, TTS applications primarily perform text and linguistic analysis to generate an audio waveform. During the text analysis, the text is analyzed and converted into full words and sentences. Any abbreviations are expanded, and expressions are identified. Linguistic analysis is then performed to understand the grammatical structure and to refine the synthetic voice's pitch and duration. The results are used to produce a spectrogram (i.e., a visual representation of the sound over a period of time) that is later converted into human-like speech.[10]

[9] Relativity, *Intelligent Voice Discovery*, https://www.relativity.com/data-solutions/customizations/app-hub/intelligent-voice-discovery.

[10] Mikiko Bazeley, *An Easy Introduction to Speech AI*, Nvidia Developer, https://developer.nvidia.com/blog/an-easy-introduction-to-speech-ai.

III. How AI is Used in Legal Practice

A. Electronic Discovery ("eDiscovery")

1. Clustering

Clustering is a form of unsupervised machine learning that involves analyzing the eligible documents in a collection (i.e., those with sentences), and building an index of "concepts" based on the words used in the sentences. Once the index is created, it can be used to search for documents that discuss similar topics and can also group documents containing similar concepts into clusters. Conceptual clusters can be used to quickly gain a better understanding of the makeup of a collection of documents and make broad coding decisions, such as denoting specific clusters as potentially responsive or potentially not responsive. Such automatic determinations should be confirmed by a human through quality control.

2. Technology-Assisted Review

AI has been used in eDiscovery for over a decade to search, classify and code documents. This is commonly referred to as "Technology-Assisted Review," or "TAR." TAR, in all its forms, uses either supervised machine learning and/or natural language processing ("NLP") to analyse textual content. NLP can be used for tasks such as machine translation, speech recognition, sentiment analysis, and topic segmentation. [11] Accordingly, TAR can only be applied to certain types of records—those that contain sufficient textual content—such as emails, letters, contracts, and some spreadsheets, but usually not images.

There are generally considered to be two types of TAR: TAR 1.0 (a/k/a "Simple Active Learning") and TAR 2.0 ("Continuous Active Learning"). Both TAR 1.0 and TAR 2.0 rely on supervised machine learning, with human subject matter experts ("SMEs") training the model to identify responsive and non-responsive documents. The classifications made by the SMEs on a subset of documents are used to construct a model that can then be applied to all documents in the collection, categorizing or ranking them by their responsiveness.

In a TAR 1.0 workflow, the model is trained using a fixed set of documents (i.e., the "training set.") Once "stabilitzation" is reached (i.e., the point at which it is determined that the model will no longer improve with a proportionate amount of effort), training is stopped, and the model is used to separate the documents into presumptively responsive and presumptively non-responsive documents. In a TAR 2.0 workflow, there is no fixed training set. Documents continue to be reviewed and fed back into the model for further training until "marginal precision" is achieved

[11]

(i.e., the number of responsive documents being found no longer justifies the effort necessary to find them).

TAR is most often used to expedite determinations of responsiveness, both in defensive and offensive review workflows. Practitioners typically use TAR to prioritize records for review, allowing them to avoid reviewing large numbers of unresponsive documents. TAR can also be used classify responsive documents into specific issues, Moreover, with proper quality conrols and in combination with other search methods, such as search terms specifying lawyer or law-firm names, TAR can be used to identify privileged documents. Indeed, regardless of the purpose for using TAR, quality control measures, such as elusion sampling of unreviewed documents, are necessary to ensure the defensibility of the review process.

Clustering and TAR are not mutually exclusive; many review workflows use both unsupervised and supervised AI techniques together to achieve more efficient results.

B. Identification and Redaction of Personally Identifying Information and Personal Health Information

AI technologies can be used to automatically identify personal information, including personally identifying information ("PII") and personal health information ("PHI"). This can be particularly useful when documents are produced in litigation or regulatory matters, to comply with privacy regulations, and when responding to a data breach.

PII and PHI detection is facilitated by a collection of AI and machine-learning algorithms that perform entity extraction, an information extraction technique that identifies key elements from text (e.g., names, addresses, phone numbers, organizations, etc.) and classifies them into predefined categories. Different methods of detection may be used, including pattern matching, NLP, and large language models ("LLMs").

C. Sentiment Analysis

Sentiment analysis uses NLP to analyze the content of text and speech, and to assess and categorize the emotional tone of that content, such as positive, negative, or neutral. This technique is particularly useful in cases involving sexual harassment, workplace investigations, or insider-trading investigations, where records that are particularly positive or negative in tone can quickly be identified and prioritized for review.

Sentiment analysis is also a valuable business-intelligence tool that provides organizations with external and back-office insight. This information allows organizations to triage and improve customer service requests and responses, monitor brand sentiment and conduct market research through social media and online sources, and track marketing campaign performance. Organizations can also use sentiment analysis internally to gauge employee engagement and

satisfaction and gain insight into employee responses to workflow and process improvements and business initiatives.

D. Language Detection and Translation

Language detection and translation use NLP or LLMs to analyze the content of text and speech to identify the primary and other languages used in those documents. AI translation services can translate text into many different languages.

In eDiscovery, language detection can be used to identify and categorize documents by language so that documents containing a specific language can be directed to reviewers fluent in that language. On-the-fly translation can be used so that reviewers who are not fluent in a particular language can translate the documents and review them, funneling only those that need it for more expensive human translation.

Organizations can also use language detection and translation to improve user experience. Some examples include translation of website or social-media content to another language based on the source country location of the website or social-media traffic, triage of foreign-language customer-service requests to specific representatives who can read and converse in that foreign language, and translation of foreign business documents.

E. Transforming Audio to Text

Audio-to-text conversion tools transcribe audible speech into text. Once transcribed, other techniques (such as clustering and TAR) can be applied to classify and review the information. While audio can be transformed to text manually or by using traditional transcription software, the use of AI for transcriptions has reduced the cost and increased the quality, and speed of audio-to-text conversions.

Audio-to-text conversion tools are also available as standalone products, but some review software platforms incorporate additional features for audio-to-text conversion.

In the context of document review, selection of review software that can provide high-quality audio transcription within the review platform can streamline the workflow, resulting in cost efficiencies for the processing and review stages of an eDiscovery project.

F. Image Classification

Image classification is a process that employs AI to identify properties within a visual image, such as a photograph, and assigns labels to represent those properties. For example, a photograph may contain people, buildings, vehicles, groundwork, etc. The extent and specificity of the labelling is dependent on the user's need(s) and the specific implementation used. A single image will typically

contain several labels. Some implementations will include the confidence of the label (i.e., a measure of the AI system's confidence that its label is accurate).

Once the labels have been created, they can be used to group images together for review and analysis. The labels can also be used to identify PII for regulatory requirements or for withholding or redacting prior to production and as an aide for people with visual impairments so they can understand the content of images they cannot see.

Review and classification of images is a key aspect for analysis and review of data from mobile phones, other portable devices, chat platform data, and social-media accounts, which typically contain many images.

G. Data-Breach Response

AI can play a significant role in identifying PII and PHI during the reporting process in response to a data breach. This is particularly important because quick and accurate identification of the information that has been compromised is crucial for determining the appropriate and response measures, including who needs to be notified and what regulatory filings may be required.

AI entity identification can be used to identify the specific pieces of PII or PHI that were involved in the breach. This could include names, email addresses, social security numbers, financial information, diagnoses, treatment(s), or other sensitive data. In addition, AI can be highly effective to link named entities to people. This process involves training an AI model to understand and recognize specific categories of information, such as people's names, within a document. The system does this by learning from a large set of training data, which includes various types of documents where the named entities and their corresponding categories have been annotated. Over time, the AI model learns to identify patterns and contexts that indicate a particular string of text is a named entity related to a particular person. For example, the AI system can be trained to recognize the difference between "Calvin Klein," the person, versus "Calvin Klein," the corporation.

Furthermore, the AI can also learn to discern contextual clues to link named entities to specific people. For example, if a document mentions that "John Doe's birthday is June 1, 1981," the AI system can determine that "June 1, 1981" is a date, is personal information (a birth date), and is associated with "John Doe" (a person). Moreover, the AI system can also use co-reference resolution techniques to link different mentions of the same person within a document. For instance, if a document initially refers to "President John Doe" and later mentions "the President," the AI system can learn that both refer to the same individual.

Importantly, while AI can be a powerful tool in identifying personal information and the associated individuals in response to a data breach, it is not infallible. Human oversight and expertise are still necessary to ensure accuracy and compliance with all relevant laws and regulations.

H. Information Governance

Recordkeeping practices evolve with changes in society. AI technology has enabled records managers and information governance ("IG") professionals to manage record systems more effectively.[12]

The first stage of any IG implementation is data classification. Without accurate and appropriate classifications, IG functions cannot work. Traditionally, classification was a manual process; a person, typically the record creator or a records manager, would review a record to determine its type (e.g., letter, memo, contract, etc.). Additional information, such as its date, author and recipients, department, or case, would also be classified.

AI systems are now available to automatically classify records, using a combination of predefined taxonomies and machine-learning techniques to parse the information within a record and apply one or many classifications. These systems continuously improve accuracy and efficiency through user interaction and feedback.

With the help of AI, search engines can understand natural-language queries and find relevant documents based on their content. AI can also enhance the searchability and organization of documents by extracting significant metadata such as dates, authors, titles, and keywords. This is especially beneficial when dealing with large records databases, as manually searching for specific information can be time-consuming.

AI can help with data governance and compliance, ensuring regulations like The Personal Information Protection and Electronic Documents Act ("PIPEDA"), Canada's federal privacy law for private sector organizations, and Quebec's Law 25, the latest and most significant privacy legislative development in Canada, are followed. AI can automate tasks related to data privacy, compliance, and security.

I. Due Diligence and Contract Analytics

AI has been integrated into several specialized due diligence review platforms. These systems use AI to identify and extract specific information, such as the parties to a lease, the location of the property at issue, and the length and price of the rental. They are also used to identify and extract particular contract clauses, such as indemnity clauses, termination clauses, or non-compete clauses. These systems can scan through large volumes of contracts and pull out the relevant sections for review. AI is also used to identify areas of potential risk in contract clauses, such as overly broad indemnification clauses or unusual termination provisions. The system can flag these areas for closer attorney review.

[12] James Lappin, *Records Management Before and After the AI Revolution*, Thinking Records, 2020

Such AI features can also compare clauses across multiple contracts to ensure consistency. They can identify deviations from an organization's standard language, which can be particularly helpful in large-scale contract reviews.

With AI becoming the norm in many areas of legal practice, it is not surprising that contract analytics has become top of mind for many corporate lawyers. Both predictive and generative AI ("GenAI") can be used by lawyers to quickly review and analyze contracts, thereby achieving efficiencies and cost savings for their clients.

Similar to due diligence tools, these tools use predictive AI to pre-train models to identify similar provisions in large contract repositories (for post-execution review). The models are typically trained on a large volume of common contractual clauses, which may reduce the amount of time required for human review of specific portions of a contract. In addition to pre-trained models, many of these tools allow users to train models on clauses or concepts that are uncommon or bespoke. Using training examples, the contract analytics tool can then identify similar clauses across a large volume of contracts. AI tools that permit users to train their own models are also generally language agnostic and can be trained in any language. Contract analysis tools can also be used to identify large groups of similar contracts or clauses, as well as anomalies and outliers. In addition to predictive AI, GenAI technologies are now being used to automate and generate clause summaries.

Emerging GenAI technologies are also being used to draft and negotiate contracts. As GenAI tools have begun to mature, we are now seeing an increased ability to identify areas where a contract differs from specific agreed terms and/or provisions that are not in compliance with the pre-defined policies or rules for certain types of agreements.

Users of these tools should be cautioned that the systems have limitations and pose risks, such as drafting inaccuracies. It will undoubtedly become more important for lawyers to adopt these technologies to remain competitive, but they should approach such new technologies with caution to ensure they are acting ethically and responsibly.

J. Compliance and Fraud Detection

AI and machine-learning algorithms can analyze patterns and trends across massive numbers of financial transactions. They can identify unusual transactions or behaviors that could indicate fraud, such as sudden large withdrawals, repeated transactions within a short period of time, or transactions in high-risk locations.

AI can also be used to automate the process of checking transactions and customer behavior against regulatory requirements. This includes monitoring for money laundering, confirming customer identities, and ensuring that all necessary information is collected and stored correctly. The risk associated with transactions, customers, or products can be automatically assessed by analyzing

historical data, customer behavior, and market trends to predict potential risks and suggest actions to mitigate such risks.

K. Legal Research

According to a study by the American Bar Association, the average lawyer spends between 16% and 25% of their time performing legal research.[13] AI has recently been incorporated into online legal research platforms to make the process more streamlined, efficient, and accurate. A study carried out by the National Legal Research Group[14] found that using AI-enabled legal research tools made the process 24.5% more efficient, saving between 132 and 200 hours of legal research per year. The researchers also found that the search results were 21% more on point that when using traditional, keyword-based search techniques.

Using NLP, legal research systems can learn the context and intent behind a user's search query, as opposed to merely matching specific words. This allows the search system to find responsive information that may contain words or phrases that are different from what the user entered, but which carry the same contextual meaning. Moreover, AI systems can extract key facts, arguments, and holdings from sources as well as summarize main points. Based on a user's prior searches and results, AI can also recommend other lines of inquiry to enhance the user's legal reasoning.

L. Legal Drafting

Through machine learning and NLP, AI can automate and streamline the legal drafting process, making it more efficient, streamlined, and comprehensible.

As described earlier, expert systems are based on a set of rules and heuristics that emulate the decision-making ability of a human expert. In legal document drafting, an expert system can be programmed with a set of rules related to the structure and content of specific legal documents. When provided with the necessary inputs, the system can generate a draft document following those rules.

Predictive text technology can also be used to speed up the document-drafting process. In this case, AI system learns from previously written legal documents and suggests the next word or phrase as the lawyer types. This can significantly accelerate the drafting process and ensure the use of standard legal language. When trained on a large corpus of legal documents, this technology can

[13] American Bar Association, 2017 Legal Technology Survey Report, Vol. V: Online Research; Steven A. Lastres, Rebooting Legal Research in a Digital Age (2013), http://www.lxisnexis.com/documents/pdf/20130806061418_large.pdf

[14] National Legal Research Group, 2018 The real impact of using artificial intelligence in legal research, https://www.lawnext.com/wp-content/uploads/2018/09/The-Real-Impact-of-Using-Artificial-Intelligence-in-Legal-Research-FINAL2.pdf

suggest relevant clauses or sections to be included in a new document based on the type of document being drafted and the specific details of the matter.

The most revolutionary use of AI, as it pertains to legal-document drafting, involves GenAI. GenAI uses deep learning and NLP to generate new content in response to a user prompt. It can be employed to automate the creation of legal paperwork, such as motions, contracts, wills, and letters, and can also be used to customize legal documents. By learning from previous documents and the preferences of the user, GenAI can suggest changes or additions to make the document more suited to the user's needs. This can include suggesting alternative phrasing for clauses, adding or removing clauses based on the specific situation, or even providing explanations for certain legal terms in plain English.

Gen AI can also help maintain consistency across multiple documents by ensuring that the same language and terminology are applied. This can be particularly useful in large-scale legal cases that involve numerous documents.

While GenAI can make the drafting of legal documents more efficient, it cannot replace the oversight and expertise of a legal professional. Lawyers must review the AI-generated content to verify its accuracy and compliance with all relevant laws and regulations, particularly in light of the fact that GenAI is prone to "hallucinations" or making up content.[15] Moreover, the use of GenAI in legal document drafting can raise ethical and legal considerations, particularly around intellectual property, liability, and confidentiality, which must be carefully considered, and where appropriate, discussed with the client.

M. Legal Analytics

AI can be used to predict legal case outcomes, a practice known as "predictive legal analytics" or just "legal analytics." Using machine-learning algorithms and NLP, AI can analyze historical case data, judicial data, attorney data to draw insights that can be used to predict the likely outcome of a future case.

Legal analytics predicts the outcome of litigation based on various inputs such as counsels' and judges' names, the jurisdiction, case type, facts, claims, and other information. Parties can use legal analytics to prepare for trial (or other stages of litigation) with a better understanding of their case's strengths, weaknesses, and overall probability of success. This information can be used to inform case strategy, including decisions whether to pursue settlement and what a favorable settlement might be. When working with limited resources, legal analytics enables parties to focus on the tasks that are most likely to advance their case objectives.

Legal analytics rely on data collected from prior cases. Using machine-learning algorithms, the AI system analyzes this data to identify patterns and correlations. For example, it might find that

[15] Cite to Mata vs. Avianca and other examples

certain arguments are more likely to be successful in particular types of cases, or that certain judges are more likely to cite opinions written by certain judges rather than others, or rule in a particular way given a certain set of facts. Based on these patterns and correlations, the AI system can predict the likely outcome of a new case. Moreover, as more case data becomes available, the AI system continuously updates its model, making its predictions more accurate over time.

Judicial analytics focus on the behavior and decisional trends of individual judges. Armed with a better understanding of how a court is likely to rule, a party can pursue arguments that are move likely to be determined in their favour. For example, if a judge consistently denies motions to compel disclosure of TAR methodology absent a demonstrated material deficiency in the producing party's production, a requesting party might elect to spend more time conferring with opposing counsel to reach a negotiated resolution, rather than engaging in costly motion practice they are more likely than not to lose.

Lawyer analytics examine metrics related to individual lawyers, like success rates and experience with different types of cases. For example, if opposing counsel has never tried a patent-related obviousness case, that weakness may inform arguments and strategy in a patent litigation. Likewise, if opposing counsel consistently avoids trials or consistently loses cases that go to trial, they may be more open to settlement discussions. Beyond individual lawyers, law firm analytics can highlight a firm's experience (or lack of experience) with certain practice areas and may help prospective clients choose the best or most experienced representation for their case.

While AI can identify trends, it cannot necessarily explain them. Nor can it account for every possible input, such as a litigator's performance on a particular day. For example, AI cannot account for certain "soft" factors like characteristics or experiences of jurors that are not reflected on questionnaires or exposed through voir dire. Certain inputs may also be confidential and thereby elude a model that is trained only on publicly available data. It is important to note that these predictions are based on statistical patterns and should not replace human judgment and expertise. Lawyers still need to interpret the results and consider other factors that may impact the particular case outcome.

N. Employment and Human Resources

Employment and Human Resources cover a wide range of legal issues, which include discrimination, harassment, wrongful termination, employment contracts, workplace safety, employee benefits, leave issues, and privacy rights, among others. AI systems can also assist HR departments in several other ways, such as compliance, recruiting, negotiations, and employee issues.

AI can be used to monitor and ensure compliance with various employment laws. For instance, it can help track employee working hours to ensure compliance with labor laws or monitor communications for potential harassment or discrimination. AI tools can also help automate parts

of the recruitment and hiring process, such as screening resumes or scheduling interviews. However, when used for such purposes, it is essential to ensure that these tools are used in ways that comply with applicable anti-discrimination laws.

AI-powered platforms can facilitate negotiations between parties, helping them find mutually agreeable solutions without the need for court intervention. These platforms use algorithms to propose solutions based on the parties' preferences and priorities. Finally, AI can analyze employee data to predict potential problems, such as employee turnover or performance issues. This can allow HR and management to proactively address these challenges before they become bigger problems.

O. Legal Spend and Legal Operations Analytics

The quality and effectiveness of legal services can be greatly improved by incorporating business principles and technology into legal operations[16] to maximize value. This transformation centers on the application of legal operations analytics,[17] which involve collecting and analyzing data to improve decision-making accuracy and facilitate performance evaluation. Strategic control over legal expenditures[18] is important to achieve this goal, with careful oversight and comprehensive reporting of all financial outlays related to legal proceedings.

Legal operations analytics can provide important insights into past litigation data, enabling legal teams to assess trends over time, the likelihood of success, anticipated costs, and potential settlement options. Legal teams can also streamline file and cost management by automating invoice review and approval processes, identifying billing discrepancies, ensuring adherence to billing guidelines, and helping in budget forecasting for legal projects. This can lead to better cost containment, reduced billing errors, and more accurate financial projections.

Law firms can also use AI-powered analytics to study client data, helping them to recognize patterns in client behaviour, preferences, and requirements. Such analysis can improve client satisfaction and retention by matching the appropriate level and type of legal services with client expectations. Additionally, AI can provide insights into market trends and competitive intelligence, enabling law firms to make informed business decisions.

AI-powered chatbots can also help lawyers to satisfy their compliance and ethical requirements. Chatbots can quickly and accurately answer compliance and ethics questions, offering reliable information and reducing the risk of non-compliance and unethical behaviour. Chatbots can also automate repetitive tasks like regulatory updates and conflict checks, allowing legal professionals to concentrate on more complex or important duties.

---

[16] *What is Legal Operations*, Association of Corporate Counsel
[17] *What is Legal Analytics and How Can Legal Ops Use Them*, Global, 2022
[18] *What is Legal Spend Management*, Thomson Reuters, 2022

P. Predictive Policing and Risk Assessment in the Criminal Justice System

AI is increasingly being used in the criminal justice system for risk and recidivism assessment, particularly in predicting the likelihood of reoffending or failing to appear in court or determining the level of threat posed by a particular individual. Likewise, predictive policing algorithms use historical crime data to predict where, when, and what types of crimes are likely to occur. These AI systems are generally designed to exclude factors like race, gender, and socioeconomic status, thereby reducing potential bias in risk assessments, but these efforts are not universally successful.

Predictive policing algorithms use historical crime data to predict where, when, and what type of crimes are likely to occur.

The use of AI in these contexts is controversial. Critics argue that these systems can reinforce existing biases in the criminal justice system, as they are often trained on data that may reflect systemic prejudice.[19] Concern have also been expressed about transparency, since the algorithms used are often proprietary and therefore not open to public scrutiny.

Q. Facial Recognition and Other Biometrics

Facial recognition and other biometrics are increasingly being used in legal cases due to their potential to provide purportedly objective and reliable evidence. However, their use also raises important legal and ethical considerations.

"Facial Recognition" is a type of biometric technology that can identify or verify a person by comparing and analyzing patterns based on the person's facial features. In legal cases, facial recognition can be used for various purposes, such as identifying suspects in criminal investigations, verifying identities in immigration cases, or finding missing persons. However, facial recognition has been criticized for its potential for misuse, bias, and infringement on personal privacy rights. There are also concerns about accuracy when used to identify individuals from certain racial and ethnic groups that may be under- or overrepresented in the training data. Several cases in the US have involved the misidentification and arrests of persons of colour for crimes they did not commit.[20]

Biometrics is the measurement and statistical analysis of a person's unique physical or behavioral characteristics. This can include fingerprints, voice patterns, iris scans, and more. In legal cases, biometric data can provide powerful evidence due to its unique and individual nature. For example, fingerprint or DNA evidence can play a crucial role in criminal cases. However, as with facial recognition, the use of biometrics raises concerns about privacy, consent, and data security.

19

20

The use and accumulation of facial recognition and biometrics data by law enforcement agencies in Canada has been the subject of debate. Some police forces have used facial recognition in criminal investigations. The Royal Canadian Mounted Police ("RCMP") maintains a national database of fingerprints and criminal records, known as the Canadian Police Information Centre. The Canadian Border Services Agency also uses biometric data, such as fingerprints and facial scans, as part of its identity verification process for travelers. Concerns have been raised about the potential for racial bias in facial-recognition algorithms, as well as the potential for the accumulation of this data to infringe on individuals' privacy rights.

In Canada, the use of facial recognition and biometrics in legal cases is regulated by various federal and provincial laws. The Privacy Act[21] at the federal level and various provincial privacy laws[22] govern the collection, use, and disclosure of personal information by government organizations. These laws require that individuals be informed of and have given consent for the collection of their personal information, which includes biometric data. Given the complex ethical and legal issues associated with these technologies, their use in legal cases in Canada is an evolving area of law and policy that lawyers should monitor.

21

22

IV. BENEFITS OF AI

The use of AI has many obvious benefits for the legal industry, most notably greater efficiency, reduced costs, and improved quality and consistency.

A. Better Quality and Greater Consistency

Data quality and consistency are critical in the legal domain, where data-driven insights and decisions can significantly influence case outcomes, client relations, and compliance with rules and regulations. Consider the use of AI for legal research. Lawyers access and search vast databases to research statutes, case law, and regulatory texts, to gather information essential for building cases, advising clients, and making informed legal decisions. If the data being searched is riddled with inaccuracies or inconsistencies, the results can lead to misguided legal strategies, misinterpretation of laws, and, ultimately, unfavourable outcomes in proceedings. For an AI system to be effectively used for research, maintaining data integrity is critical to ensuring quality and consistency in the results. This is often easier to obtain using computer-assisted processes than purely manual ones. This has been shown to be the case particularly in the context of TAR versus manual review for document productions.[23]

B. Increased Defensibility

When an AI system is defensible, it means its decisions are transparent, interpretable, and justifiable. This builds trust among lawyers, clients, and the judicial system, thereby enhancing the credibility of both the AI tool and the legal practitioners using it. While not all AI systems are transparent, many AI methods can be readily subjected to objective validation processes to demonstrate their validity, and the reliability of their outputs. This is particular true as to the use of TAR in eDiscovery, where the calculation of metrics such as recall[24] and precision[25] are common.

C. Greater Efficiency

The use of AI systems facilitates the effective use of resources by automating tasks and providing accurate insights, allowing legal professionals to focus on more complex, value-added aspects of their work. Through the automation of repetitive manual and often time-intensive processes and tasks, AI systems offer practitioners an invaluable commodity: time. This efficiency can contribute meaningfully to the overall productivity and profitability of a legal practice.

---

[23] See, e.g., Add cite to Grossman and Cormack, Richmond Journal (2011).
[24] Cite to Grossman-Cormack TAR Glossary for definition.
[25] Cite to Grossman-Cormack TAR Glossary for definition.

D. Permits Lawyers to Focus on Higher-Level Work

AI systems, through their demonstrated validity and reliability, can streamline various repetitive, time-consuming tasks that lawyers have traditionally performed, such as document review, due diligence, contract analysis, proofreading, and legal research. As a result, efficiencies are gained that allow lawyers to redirect their efforts from these time-consuming tasks to high-level, intellectual work requiring judgment that adds more value to clients and the legal practice by elevating the level of service provided and ultimately saving on costs.

E. Cost Savings

Through automation, legal professionals can reduce the cost of previously highly manual tasks and focus on more complex, high-value work, thereby increasing overall productivity and reducing the hours billed for time-consuming and primarily ministerial tasks. For example, document reviews that previously required large teams of contract reviewers and many months to complete can now be done in just a couple of weeks, by a smaller team, by leveraging AI-based TAR tools. In this scenario, the cost savings can be passed on to the client, and the lawyer's time is freed up to focus on quality control and other work important to the case. Furthermore, AI systems are less prone to errors compared to humans, particularly in tasks that involve large datasets or repetitive processes, such as document review. By minimizing human error, AI helps avoid potential costs associated with fixing mistakes and mitigating any legal issues that arise from such mistakes.

F. Potential Increases in Access to Justice

Embracing AI within the legal domain holds substantial promise for amplifying access to justice for a wider group of individuals. Streamlining various aspects of civil litigation, for drafting of court filings, opens the door to the justice systems for a larger and more diverse population. Furthermore, legal practitioners, once weighed down by heavy caseloads, gain the capacity to manage additional cases, thereby extending essential legal assistance to more clients. Similarly, smaller law firms, that once struggled to compete with larger ones, now have access to the same advanced legal research tools and analytics, levelling the playing field and enhancing the quality of legal representation.

V. CONSIDERATIONS WHEN USING AI

While AI presents tremendous opportunities for productivity and innovation, its practical application requires a solid understanding its limitations and challenges. As AI tools become intertwined with the fabric of legal practice, understanding their intricacies is both beneficial and essential. Some of these challenges and considerations that apply to legal applications are discussed in this section.

A. Data Quality: Not All Data is Created Equal

To mitigate issues such as bias or invalid predictions, it is essential to use comprehensive and representative datasets for training AI systems. Data imputation techniques can also be used to estimate missing values in the dataset. However, these techniques have their own limitations and must be applied carefully. Ultimately, the quality and completeness of the data used to train AI systems are critical determinants of their performance and fairness.

At its core, AI is based on data. Therefore, data quality and completeness is crucial for accurate and reliable results. Incomplete data can lead to inaccuracy, bias, and/or overfitting.[26] Incomplete data can also lead to legal and ethical issues, particularly if an AI system's decisions or predictions negatively impact certain individuals or groups due to gaps or errors in the data.

B. Correlation vs. Causation: Seeing Patterns vs. Understanding Them

When diving into data analysis, particularly with AI, it is helpful to understand the difference between two foundational concepts: correlation and causation.

Correlation refers to an observed relationship or connection between two or more variables. When one variable changes, there is a consistent, observable pattern in the other, in the same direction. For example, a rise in online searches for winter coats might correlate with colder months. However, the correlation does not necessarily mean that colder months caused the increase in searches.

Causation implies a cause-and-effect relationship between two or more variables. Using the previous example, if causation is shown, then we can assume that online searches for winter coats will increase because of the presence of colder months. In the context of AI, identifying causal relationships can be challenging, as AI models are typically trained on correlations between variables.

[26] Overfitting occurs where an AI model learns to perform well on its training data but performs poorly on new, unseen data. This is because the model has effectively memorized the training data, including its gaps and inconsistencies, rather than learning to properly generalize from it.

Understanding the difference between correlation and causation is essential for interpreting AI outputs. Misinterpreting correlation as causation can lead to flawed decision-making and inaccurate predictions. Hence, human oversight and domain knowledge are crucial in interpreting the results generated by AI systems.

C. Bias: Unwanted Baggage in AI

Bias in AI refers to systematic errors or prejudices in the outputs produced by AI systems. It can be categorized into two types: intentional and unintentional bias.

Intentional bias occurs when an AI system is deliberately designed to favour certain outcomes over others. For example, a loan-approval algorithm might be intentionally biased to favour certain types of applicants (e.g., those with higher credit scores) if it has been specifically programmed to do so. While intentional bias can sometimes be used for legitimate purposes, such as promoting fairness or diversity, if misused, it can also lead to discriminatory outcomes.

Unintentional bias occurs when an AI system produces biased or unfair results due to issues in the training data or errors in the design of the algorithm. For example, suppose a facial recognition system is trained predominantly on images of light-skinned individuals. In that case, it may perform poorly on individuals with darker skin tones.

Bias in AI can significantly affect the outcomes and predictions of a system. It is therefore essential to take steps to identify and mitigate bias in AI, such as using representative and diverse training data, regularly auditing and testing AI systems for bias, and incorporating fairness metrics into the design of AI algorithms. Recognizing and addressing these biases is important, especially in a legal context where equity and impartiality are core principles.

D. Equitable Access and Other Fairness Considerations

AI systems have the potential to help close the access-to-justice gap, yet, at the same time, there is also a fear that AI will increase inequities, favouring of those who can afford the benefits it provides. This could further exacerbate existing disparities, leaving marginalized communities behind.[27] A concerted effort is needed to ensure a two-tiered system is avoided and that developers strive to close the gap in Canada's access to justice problem.

Fairness in AI systems refers to the absence of discrimination or favouritism towards any individual or group based on protected characteristics such as race, gender, age or religion.[28] For

[27] See Abdi Aidid & Benjamin Alarie, *The Legal Singularity: How Artificial Intelligence Can Make Law Radically Better*. p.146-150; Drew Simshaw. *Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services*. 24 Yale J.L. & Tech. 150 (year?)

[28] Dwork et at., 2012; Ferrara E. (2023) Fairness And Bias in Artificial Intelligence: A Brief Survey of Sources, Impacts and Mitigation Strategies. https://arxiv.org/abs/2304.07683

instance, an AI tool used to predict the likelihood of recidivism should not disproportionately categorize persons of colour as high-risk due to biases in the training data or the algorithm.

Equitable access to AI in the justice system means that datasets close societal gaps and minimize inequity by making sure that AI models are trained on appropriate and representative data that provide relevant, accurate, and unbiased outputs.

E. Defensibility and Validation: Ensuring AI's Credibility, Consistency and Safety

Decisions based on AI systems must be defensible to comply with legal principles, industry standards, and ethical guidelines. Defensibility concerns the ability to justify and defend the decisions made by the AI system. This involves ensuring that, whenever possible, the AI system is built and trained in a transparent, interpretable, and explainable way. This is particularly important in a legal context where the decisions and actions taken based on AI recommendations can have significant legal implications.

When transparency is not possible, independent validation becomes even more important. Validation refers to an objective assessment of whether the AI is working as intended (i.e., valid) and produces accurate results under substantially similar circumstances (i.e., consistency and reliability). This typically involves testing the system on a separate validation dataset not used in the training process to evaluate the system's performance and generalizability to new, unseen data. Validation should also consider other aspects like fairness (i.e., the system should not be biased against certain groups), robustness (i.e., the system should perform well even under non-ideal conditions or when it is subject to intentional efforts to have it malfunction), and safety (i.e., the system should not cause harm or undesired outcomes).

F. Opaqueness

Opaqueness refers to the lack of transparency or clarity in an AI system's predictions or decisions, and AI systems that are opaque are often referred to as "black box." AI systems, particularly those based on complex machine learning algorithms like deep learning, often involve complicated computations and large amounts of data. Testing is almost always needed to ensure the reliability and credibility of the system's outputs, particularly when the cannot be explained. For example, if an AI system is used to recommend a patient's treatment in healthcare, doctors and patients would want to understand why that recommendation was made and that the prediction is highly accurate. Similarly, in the legal or financial sectors, explanations may be required for decisions that have significant legal implications in addition to proof of accuracy.

Efforts to address the opacity problem in AI are often focused on developing techniques for "explainable AI" ("XAI") or "interpretable AI." These are AI systems that are not only capable of

making decisions or predictions, but also of providing explanations for those decisions that are understandable to humans. This is a growing field in computer science.

G. Accountability: Ensuring AI Operates Responsibly and Ethically

As AI technology rapidly advances, it raises significant ethical and accountability concerns that must be addressed. AI currently lacks a fundamental quality that humans possess: judgement. This means that:

1. AI systems should be designed and used in a way that treats all individuals and groups fairly. It should not discriminate against, or harm certain groups based on legally protected characteristics;
2. AI systems should operate transparently to the degree possible. Their decision-making processes should be explainable and understandable by humans, at least at some level of abstraction;
3. There must be clear accountability for the decisions made by AI systems. This involves establishing a specific person or entity that is responsible when an AI system makes a mistake or causes harm. It also involves creating mechanisms for auditing and overseeing the use of AI;
4. AI systems must comply with relevant laws and regulations regarding data protection and they must use data in a way that respects individuals' privacy rights; and
5. AI systems should be designed and used in a way that ensures the safety and security of individuals and society.

H. Privacy and Security

Protection of personal and confidential information and prevention of unauthorized access are at the forefront of AI's integration into legal practice. There is a significant concern that private or confidential information may be exposed in ways that can violate solicitor-client privilege or privacy regulations. As AI systems and models evolve, so will creative nefarious methods to hack into data and algorithms, affecting the trustworthiness of results. It is crucial to be aware of some of the threats currently at play and to remain current with respect to known privacy risks.

Many AI systems rely on large amounts of personal data to function effectively. It is essential to ensure that this data is collected with consent, used only for the purpose(s) for which it was collected, and used appropriately. Strong data governance practices and anonymization techniques can help protect individuals' personal information.

Like any other software systems, AI systems can be vulnerable to cyber attacks of all sorts. Implementing robust security measures to protect these systems from threats is important. Further, since legitimate AI technology that is widely available can be used maliciously, such as by creating

deepfakes or through adversarial attacks, such systems can be deceived through subtle modifications of input data. Accordingly, it is essential to consider security in AI design and to ensure that robust processes in place to prevent such attacks.

I. Authentication and Admissibility Issues

The field of law requires both the precision and validation of evidence and information.

Authentication refers to the process of verifying the reliability and integrity of data or information. In the context of AI, authentication might involve validating the AI algorithms, verifying the data used to train the AI, and ensuring the AI system hasn't been tampered with. Conversely, admissibility, refers to whether the evidence is allowed to be presented in court. For AI-generated evidence to be admissible, it usually needs to be relevant to the case at hand, shown to be reliable, must not be unfairly prejudicial or misleading, and must comply with legal standards.

The reliability of the AI system that generates evidence can be demonstrated by showing that the AI system is widely accepted and used in its field, that it's based on sound scientific principles, and that it produces accurate and consistent results. An example of inaccurate results is Deepfakes. Deepfakes are synthetic media in which a person in an existing image or video is replaced with someone else's likeness. These manipulations can affect legal cases or public perception. "As deepfake technology improves and it becomes harder to tell what is real, juries may start questioning the authenticity of properly admitted evidence, which in turn may have a corrosive effect on the justice system."[29] Verification tools, such as reverse image searches, are necessary countermeasures against such deceptive tactics.

Given the complexity and novelty of AI technology, there can be significant challenges in authenticating AI systems and in demonstrating the admissibility of AI-generated evidence. This evolving area of law is likely to continue developing as AI becomes increasingly used in the legal field.

J. Ethical Considerations: AI's integration into law must align with the profession's stringent ethical standards.

AI's use in law brings several ethical considerations to the forefront, including:

1. Competency: Legal practitioners must maintain a proficient understanding of the AI tools they employ, ensuring accurate and ethical use[30].

[29] Rebecca A. Delfino, *Deepfakes on Trial: A Call Trial: A Call To Expand The Trial Judge's Gatekeeping Role To Protect Legal Proceedings from Technological Fakery*, (Hastings Law Journal 2023), p.297

[30] Section 3.1-4-.4A of the LSO Rules of Professional Conduct states that, "to maintain the required level of competence, a lawyer should develop an understanding of, and ability to use, technology relevant to the nature and

2. Protecting Confidentiality and Privilege: Any AI tool or system used in legal contexts must uphold the sacred trust of client confidentiality and privileged information.
3. Supervision: Continual oversight of AI systems is required, ensuring they align with legal and ethical standards.
4. Quality of Legal Services: While AI can help automate certain legal tasks, it's essential to ensure that this doesn't compromise the quality of legal services. AI should not replace the need for competent legal advice from a human lawyer.

These ethical issues highlight the need for careful oversight and regulation of AI in the legal field, and ongoing research and dialogue about the responsible and ethical use of AI in law.

area of the lawyer's practice and responsibilities. A lawyer should understand the benefits and risks associated with relevant technology, recognizing the lawyer's duty to protect confidential information set out in section 3.3.

VI. CURRENT AND FUTURE REGULATORY RESPONSES

Artificial Intelligence is undoubtedly the most transformative technology of our time. The global increase in the use of AI is calling for regulation. While there are no existing federal AI-specific regulations yet in Canada, the Artificial Intelligence and Data Act, which is part of proposed Canadian Privacy Legislation and Bill C-27, aims to fill this void. The use of AI is covered in Quebec's Law 25.

A. Canadian Privacy Legislation and Bill C-27

Canada's PIPEDA is a federal law that regulates the collection, use, and disclosure of personal information, and requires protection of such information against unauthorized use or theft.[31]

Bill C-27 would repeal Part I of PIPEDA and replace it with a new Consumer Privacy Protection Act, which would allow organizations to de-identify and anonymize personal information, along with requirements for transparency around the use of automated decisions systems (including systems using predictive analytics or machine learning).[32]

Some Canadian provinces have correlates to Canadian federal privacy laws. One example is Quebec Law 25, which includes requirements for reporting on the use of biometric data and reporting of security incidents.[33]

B. Artificial Intelligence and Data Act

As part of Bill C-27, Canada introduced the "Artificial Intelligence and Data Act" ("AIDA") to regulate AI in the Canadian private sector.[34]

The proposed legislation includes requirements on persons responsible for "high-impact" systems to adopt measures to assess and mitigate risk of biased output or other harms stemming from the system.

The definition of a "high-impact system" under the AIDA will be established by regulations. AIDA companion documentation provided by the government of Canada indicates that high-impact systems would be those that impact, among other things, health and safety and human rights.[35] It is yet to be determined whether the AI algorithms used in eDiscovery applications would be considered high-impact systems, though that seems unlikely.

---

[31] https://laws-lois.justice.gc.ca/ENG/ACTS/P-8.6/FullText.html

[32] https://www.parl.ca/DocumentViewer/en/44-1/bill/C-27/first-reading at CPPA ss. 62(2)(c) and 2(1) and Consequential and Related Amendments s.4.

[33] https://www.canlii.org/en/qc/laws/astat/sq-2021-c-25/latest/sq-2021-c-25.html

[34] Supra note _ at s.4.

[35] https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document

AIDA requires managers of a high-impact system to make available plain language descriptions of such system, including how the system is intended to be used; the type of information used in the system; any decisions, recommendations, or predictions that are intended to be made; and mitigation measures.

The AIDA also provides for the establishment of an Artificial Intelligence and Data Commissioner, monetary penalties and offense provisions to enforce requirements.

C. AI Regulations in other Jurisdictions

The General Data Protection Regulation 2016/679 (GDPR) is currently the most comprehensive data protection law in the world. It came into effect May 25, 2018. The GDPR repealed the Data Protection Directive 95/46/EC, which was adopted to deal with the rise of the internet.[36] The GDPR protects the fundamental rights and freedoms of natural persons, their right to protection of personal data in balance with other fundamental human rights in the EU and international business.[37] The GDPR does not directly govern AI; however, it does address risk processes used by AI, including profiling, erasure, destruction of data, and automated individual decision making (ADM).[38]

The EU Artificial Intelligence Act ("AI Act") is the first comprehensive AI legislation that has been approved and is expected to be in full force by 2026.[39] The AI Act interacts directly with GDPR data governance, as well as with other law enforcement directives and human rights rules regarding the design, development, and use of certain high risk AI systems and also certain uses of remote biometric identification systems. The AI Act also minimize the risk of algorithmic discrimination. The legislation in its present form classifies specific levels of AI risk, (1) unacceptable risk, (2) high risk, and (3) low or minimal risk, with an emphasis on high-risk systems.[40] As of December 2023, the A.I. Act _____.[41]

The Illinois Biometric Information Privacy Act ("BIPA") is an example of a jurisdiction-specific and industry-specific regulation in the United States, of which there are a number. It is one of the most feared and challenged privacy laws after the Illinois Supreme Court unanimously held in *Rosenbach* v. *Six Flags Entm't Corp.*, that private entities cannot collect biometric data from

---

[36] EU General Data Protection Regulation 2016/679 art. 1, 94, 4(2), 22, 2016 O.J. (L 119) 1-88. (*enforceable May 25, 2018*).
[37] Article 8(1) of the European Convention on Human Rights, 213 UNTS 221, Cmnd 8969 (1953).
[38] *Supra*, note 2, art. 4 (2), 22.
[39]
[40] Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL LAYING DOWN HARMONISED RULES ON ARTIFICIAL INTELLIGENCE (ARTIFICIAL INTELLIGENCE ACT) AND AMENDING CERTAIN UNION LEGISLATIVE ACTS- COM/2021/206 final, *at* https://eurlex.europa.eu/legalcontent/EN/TXT/?qid=1623335154975&uri=CELEX%3A52021PC0206
(text of the bill, A.I. Act 1.2, 3.5, 5.2.2, 10, 29 (6); GDPR art. 22, 30 (*A.I. Act, GDPR interplay*).
[41]

consumers without their consent.[42] Ruling against Six Flags that collecting thumbprints without permission, despite no actual harm to the claimant, violated § 15 of BIPA.[43] A plethora of BIPA lawsuits have subsequently been brought and Courts seem to be protecting critical privacy interest over business interests. For example, Texas has recently enforced biometric law in two recent cases brought against Metaform Platforms, Inc., f/k/a Facebook and Google.[44] Unlike, Illinois BIPA, Texas Capture or Use of Biometric Identifier Act ("CUBI") does not have a private right of action. CUBI can only be enforced by the State.[45] The *Metaform Platforms* case focused more on lack of consent to use Biometrics AI. on customers, while the *Google* case focused more on deceptive advertising around the use of the BI collected, including that Google does profiling with A.I.[46]

New York has implemented a new law to regulate bias in the use of AI in hiring decisions. It became effective on July 5, 2023, and is known as the Automatic Employment Decision Tool ("AEDT") law.[47] The bill requires that a bias audit be conducted on any automated employment decision tool prior to use of the tool. The law also requires employees and job candidates be notified of the use of the tool during the hiring process. Violations are subject to civil penalties.[48] These acts are representative of the approach the US has generally taken, which has tended to be more localized and sector specific than the approach taken in the EU and anticipated to be taken in Canada.

A new US Regulatory bill was submitted to Congress in ___, that would require those using GenAI to include a disclaimer "this output has been generated by artificial intelligence" on all AI-generated content including photos, videos, text, audio, etc. The use of AI without this disclaimer would violate FTC §18 (a) (1) (b), 15 U.S.C. 57a (a) (1) (B), which prohibits unfair or deceptive acts.[49] It remains to be seen whether this bill will pass.

[Need to address the October 30 Executive Order]

[42] Rosenbach v Six Flags Entm"t Corp., 2019 IL 123186, 5, 6, 9. (Jan 25, 2019). (*¶ 34-35, BIPA injury is not just technical, but real and significant with no recourse for the injury*).
[43] Illinois Biometric Information Privacy Act (740 ILCS 14/1 et seq) (West 2016). *(§ 15 of BIPA-regulates the retention, collection, disclosure and destruction of biometric data*).
[44] State of Texas vs. Meta Platforms, Inc., f/K/a Facebook Inc. No. 22-0121, 20-24 (D. Tex. Filed Feb 14, 2022); Complaint, at 8-16, 49 State of Texas vs Google LLC No. ____ (D. Tex. File Oct 20, 2022). (*holdings not clear*).
[45]
[46] *Id.* at 49, *supra* note 13.
[47]
[48] New York City, N.Y., Local Law No. 144 Int. No.1894-A (West 2021).
[49] AI Disclosure Act of 2023, H.R. 3831,118th Cong. §2 (a), (b) (2023). (Generative A.I.Disclaimer).

VII. LOOKING AHEAD

A. Authentication and Admissibility Issues

Evidence has always been fundamental to justice. With the emergence of AI systems, we are seeing new types of evidence emerge, such as deepfakes, that require special attention and scrutiny. These new types of data are likely to necessitate the need for experts who can delve into the intricacies of AI-generated evidence to determine if it is authentic and suitable for admission into evidence.[50] As AI continues to advance, so do the challenges of verifying the origin and integrity of digital evidence.[51] Litigators, in particular, need to stay abreast of legal developments in this area.

B. All Manner of Deepfakes

The rapid advancement of deepfake technology poses a pressing concern. AI-generated audiovisual clips can be alarmingly convincing and capable of simulating real individuals' speech and movements. In the context of litigation, if a deepfake video can be presented as genuine evidence, it will seriously undermine the integrity of proceedings and impact court decisions.

Lawyers may soon find themselves leaning on cutting-edge tools and experts to affirm the authenticity of evidence. This technology's potential was illustrated recently when comedian Jordan Peele partnered with BuzzFeed to produce a video that seemingly had former U.S. President Barack Obama making startling comments.[52] Yet, the voice behind the revelations was Peele's, highlighting how seeing isn't necessarily believing anymore. As technology evolves, the legal world must remain vigilant, ensuring that evidence remains genuine in an era where digital impersonations are increasingly and unfortunately seamless.

To help detect fake audio and video, companies like Microsoft and Google (and others) are researching and developing ways to help detect deepfakes. One example is Microsoft, which is adding a feature to its Bing Image Generator tool to help people see if pictures or videos were made by AI.[53] This feature uses a special tag with information about the provenance of the content, commonly referred to as "watermarking." Google is developing something similar that will show if visual media was created by AI.[54]

[50] Cite to AI as Evidence and the GPT Judge.
[51] Cite to AI as Evidence and the GPT Judge.
[52]
[53]
[54] Microsoft Pledges to watermark AI-generated images and videos, Kyle Wiggers, (Tech Crunch) May 23, 2023

C. Generative AI ("GenAI")

Ongoing research and development efforts with respect to Generative AI are currently operating at a breakneck speed. GenAI relies on a combination of deep learning and NLP trained on vast datasets—typically scraped from the Internet—to generate new material from a human prompt. GenAI is quickly expanding from text (i.e., large language models) to images, audio, and video. Moving forward, we can expect this technology to become more ubiquitous and easier to use. Moreover, the quality of converting text to images, audio, or video (and converting images, audio, or video to text) will continue to improve. Meta's Audiocraft is a recent example of this technology.[55] AudioCraft consists of two primary models available to users, MusicGen and AudioGen. The MusicGen model allows users to generate music from text, whereas the AudioGen generates public sound effects from text. As described earlier, the creation of high-quality computer-generated images, audio, and video will make it more difficult to tell the difference between human- versus computer-generated content. As we have already seen, computer-generated content also raise copyright ownership issues.[56] To date, most content created by GenAI in response to human prompts has not been considered subject to copyright protection.[57]

D. Access to Justice

AI systems and solutions have the potential to help close the access-to-justice gap through the rise of "efficiencies, democratiz[ing] access to legal information, and help[ing] consumers solve their own legal problems or connect them with licensed professionals who can."[58] We can expect to see more individuals representing themselves in court using AI software to draft legal documents. AI solutions are being developed to improve various aspects of the legal process from intake to case management and judicial engagements.

E. Robotics

While AI refers to software, robotics generally involves hardware that perceives and acts in the world. Robotics and AI are widely used in many industries, such as manufacturing, automotive, packaging, and surgery. Historically, programming a robot to complete a simple human task (e.g., cleaning up a spill) required sets of complex instructions to account for all the obstacles it may encounter in the real world. Due to advances in AI, robots can now access a large amount of data to assist in its decision making instead of requiring hardcoded instructions. For example, Google's RT-2 robot incorporates large language models, improving the robots reasoning and improvisation

55

56

57

58 Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services., Drew Simshaw. (24 Yale J.L. & Tech. 150), June 20, 2022

skills.[59] By leveraging the large data available in these models, the RT-2 robot can analyze and navigate the surrounding environment.

F. General or Strong AI

IBM defines Strong AI or Artificial General Intelligence ("AGI") or General AI as "a theoretical form of AI used to describe a certain mindset of AI development. If researchers can develop Strong AI, the machine would require an intelligence equal to humans; it would have a self-aware consciousness that has the ability to solve problems, learn, and plan for the future."[60] According to most definitions, AGI would meet or exceed the capacity of humans at most if not all tasks. There is considerable debate in the field of computer science as to whether and when AGI will be achieved, with estimates ranging from several years to never.

Today, we engage primarily with narrow or weak AI systems designed for specialized tasks, such as analyzing legal documents or forecasting case outcomes. Looking ahead, we can see a time where AI becomes more advanced, where it would be equipped with cognitive expertise to tackle intellectual tasks that currently only a human can undertake. For the legal field, this paints a picture of a future where lawyers work alongside AI counterparts that can grasp the nuances and subtleties of laws and facts, and gleaning insights in seconds from massive volumes of data their AI "colleagues" processed instantaneously.

G. Superintelligence

As defined by Merriam Webster, Superintelligence refers to "an entity that surpasses humans in overall intelligence or in some particular measure of intelligence." This kind of AI is often depicted in science fiction as AI systems that greatly exceed human capabilities and often produce dystopian outcomes. Such leaps in AI capabilities would be revolutionary for every sector, including the legal world. Imagine a legal system empowered by an AI that could sift through every known legal document in moments, predict the outcome of court cases with astounding precision, and draft just and effective legislation. The rise of superintelligence raises pressing ethical questions such as, how far we let our trust in machinery extend and whether human lawyers and judges can be completely replaced by AI, or there are values that dictate that certain decisions (e.g., child custody or criminal sentencing) remain human functions.

59

60

VIII. CONCLUSION

[We need to add a Conclusion section]