

Artificial intelligence:

What is AI and
will it really
replace lawyers?

©2018 Lawyers' Professional Indemnity Company.

This article originally appeared in LAWPRO Magazine (Vol. 17 no. 1).

It is available at www.lawpro.ca/lawpromag

The practicePRO and TitlePLUS programs are provided by LAWPRO

If you scanned social media or the headlines in many online or print-based newspapers or magazines published in 2017, you were pretty much guaranteed to see posts and articles on artificial intelligence (AI).

Most of these articles suggest that AI is in the process of fundamentally changing our lives at work, home and play. And if you believe the comments in these articles, the good news is that we will have more free time to enjoy virtual-reality worlds and have our self-driving cars take us around the countryside. The bad news is that many people, including lawyers, will supposedly lose our jobs to AI technology and robots. There is no doubt, along with other major disruptions (See “Perspectives on the future of law: How the profession should respond to major disruptions” at page 5), AI technologies have and will bring changes to the legal services arena. This article attempts to sort out the hype and reality of how AI will impact the legal profession.

What is AI?

To really understand the impact AI will have on the legal profession, we should start with a clear understanding of what AI really is. This is difficult as even AI experts can't seem to agree on a definition for AI. To further complicate things, the definition of AI has changed over time as computers have become increasingly capable. For example, while some considered it AI when it was newly available, optical character recognition (the ability of a computer to recognize letters in a scanned image of a document) is now considered a routine technology by most people.

At the simplest level, you can say AI is the capability of a computer or machine to imitate intelligent human behavior. To add some details, it means a machine that can learn and think. A machine that is “smart” enough to know or recognize things and mimic human cognition for problem solving. As you will see, there are many different AI technologies involved in mimicking human senses and thinking. Let's look at them individually, and then discuss how they can work together. As you will also see, higher levels of human functionality only become possible when different types of AI work together.

Text/speech manipulation

Using skills that were learned at an early age, with little thought or effort, most people engage in many oral conversations and countless instances of reading text on a daily basis. Text and speech manipulation seems very easy to most of us. But when you break it down, there's a lot happening here. It's much harder than it looks.

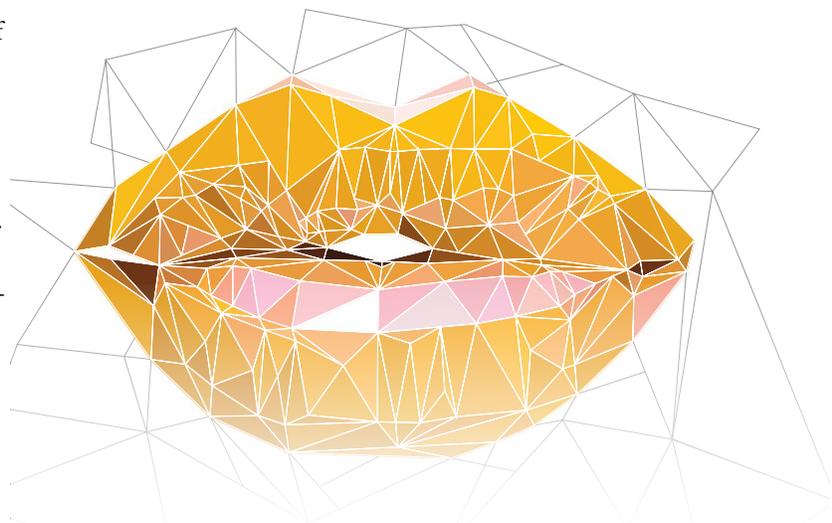
The first version of DragonDictate® software was released in 1982. Each word had to be enunciated individually with a slight pause between them so it could recognize the intended word by analyzing

the sound pattern it heard. It had no understanding of the words it was translating and it easily confused words like to, two and too or there, their and they're. The widely used current version of this software, Dragon Naturally Speaking®, is far more capable. It does an excellent job of recognizing words in a continuous stream of speech and will improve its accuracy by learning the nuances in a particular person's voice. As it converts spoken words to text on a screen it can simultaneously correct grammar and pick the correct homonym by looking at the other words in the sentence. However, it still doesn't really understand the words it is transcribing.

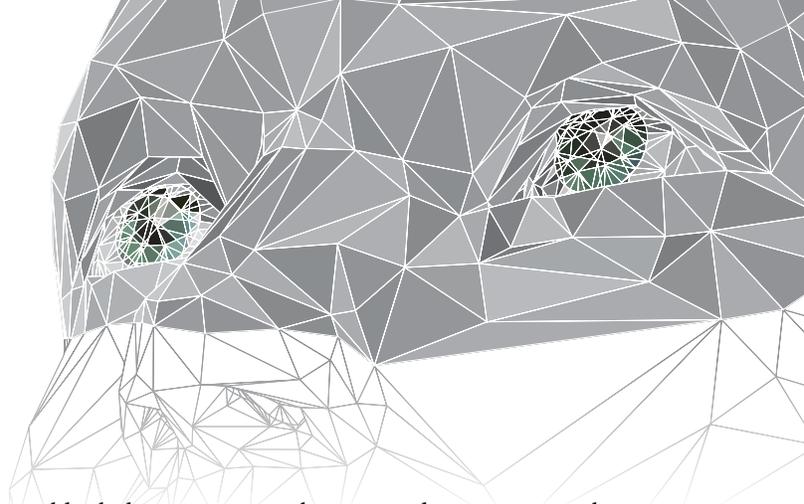
Text readers can convert text into words spoken by a very human sounding voice. While early text readers sounded robotic and were hard to understand, Google's DeepMind® AI allows computers to mimic the human voice in a manner that is virtually indistinguishable from a real human voice.¹

While it takes some effort to learn how to use Dragon Naturally Speaking (you need to learn 20 or so voice commands to use it effectively), it is a tool that can make most lawyers more efficient as it lets them put words on a screen faster than they can type. It is very helpful for answering emails or drafting documents. It is less helpful for drafting long documents with complex formatting. Voice-recognition software is being used in other situations. The basic features on most smart phones can now be operated with spoken commands, as well as hands-free smart speakers like the Amazon Echo and Google Home. In the not too distant future you will be talking and listening to your car and most of the devices and appliances in your house.

While text readers are currently used primarily by lawyers with visual impairments, they can be useful for any lawyer looking to have a document read to them for proofing or review purposes, or while they are commuting.



¹ Listen to some sample audio files and read more about how they are created here: deepmind.com/blog/wavenet-generative-model-raw-audio/



Vision

The human eye is truly amazing. It can see in light and dark and change its focus from near to far and back again in an instant. With the assistance of powerful software, computers are learning to see too, and they are also gaining some of the visual processing capabilities that humans have.

Machine vision couples a visual input with analysis and some kind of mechanical device. An example would be a device that sorts fruit by ripeness as it passes by on a conveyor belt.² But it has moved far beyond recognizing when a green tomato goes by – more advanced AI technologies are allowing computers to be smart enough to recognize people and objects in a picture. The tagging feature in Facebook® has the ability to recognize you or one of your friends in a picture you just uploaded. Google's image search has the ability to identify what a picture contains (e.g., a dog, a forest, mountains or a sunset). These technologies are not perfect as they don't accurately identify the contents or faces in a picture 100 per cent of the time, but they are getting very impressive.

Decision trees

Moving up a level we have decision trees. While vision and text manipulation may seem a bit abstract when it comes to the work that needs to get done on a daily basis in a law office, decision trees will seem more relevant as they can directly mimic the very specialized work that lawyers do.

A decision tree is a logical structure that contains every question a lawyer would normally ask when handling typical factual scenarios and legal issues for a certain type of legal matter. For example, consider building a decision tree that would do the intake on a will matter. You would first assemble the questions to gather the basic identity and background information of the client and beneficiaries. Further questions would draw out the client's instructions on the basic provisions that go in every will (e.g., executor(s) and alternate executor(s), specific bequests and gifts, gift overs, etc.). A will matter intake decision tree would consist of a few hundred or more questions in many branches. The answers to certain questions would trigger the decision to answer or skip further questions in other branches (e.g., the question asking if any beneficiaries were minors would trigger the need to ask or skip the questions on setting up trust provisions). Once all the relevant questions are asked, the answers could be dumped into a document assembly engine which could create a will that has all the relevant clauses based on the client's information and instructions in a matter of seconds.

A decision tree "thinks" in much the same way that a lawyer would, albeit in a much more organized fashion. A decision tree system

would ask the questions in the same order every time whereas an experienced wills and estates lawyer would ask all the same questions, but the order might change (unless she was using a checklist).

If you had enough time and money, you could build a decision tree that would handle every possible scenario that might be encountered on a will intake. This decision tree would likely have thousands of questions but it would be impractical as most of the questions would not be relevant to most clients. Some lawyers will argue that every matter they handle is unique and requires a custom solution which only a lawyer can provide after having done a thorough analysis. While there are extremely complex matters that are unique if you analyzed them down to very detailed factual level, in most areas of the law there are one or more common factual scenarios and legal issues that repeat themselves over and over in the majority of matters. The trick is building a decision tree that will ask all necessary questions to draft a will in the majority of situations, while at the same time flagging when manual intervention is needed because the facts or legal issues are not properly addressed in the questions within the decision tree. Typically a lawyer would still meet with the client to review the will and make sure it was correctly drafted and that the client understands the provisions in it.

Decision trees can be built into website or smart phone apps and could be helpful to automate the intake process on a wide variety of legal matters or for some stages of some types of matters (e.g., gathering the information for a financial statement on a family law matter). Spending 30 minutes with the client to review information provided from an automated intake is more efficient than spending two hours with the client transcribing all the background details of the matter. Some online forms sites are using decision tree and document assembly technology to automate the creation of forms that are provided directly to clients.

Natural language processing

Natural language processing (NLP) takes things to the next level and involves creating AI that can understand how humans understand language. There are two approaches to NLP: rules-based NLP and statistical NLP.³

² Watch this video of a machine that sorts red and green tomatoes: youtube.com/watch?v=Bur5g2rvXog

³ Paraphrased from *What is Natural Language Processing (NLP)?* by Dr. Rutu Mulkar-Mehta (ticary.com/2017/12/12/what-is-nlp.html).

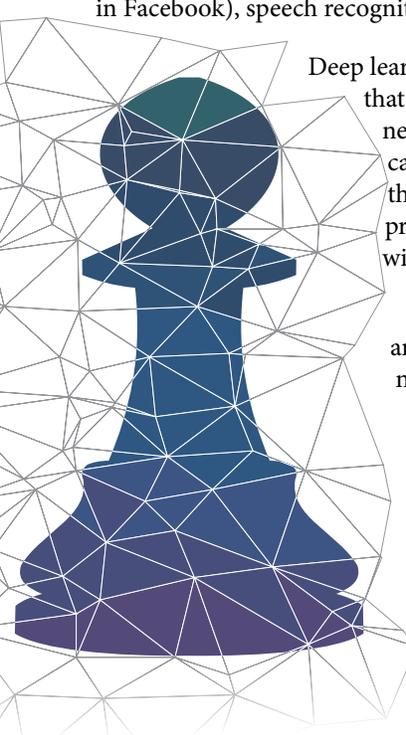
Rules-based NLP involves common sense knowledge which is all the inherent background knowledge human beings take for granted in our daily lives (e.g., freezing temperatures cause hypothermia or hot coffee will burn skin). Encoding common sense knowledge is a very manual and time-consuming process because it isn't written down clearly anywhere and it is difficult to identify all the rules required to understand something. Further, some common sense knowledge which humans inherently and easily understand can't be explained to a computer with a simple rule (e.g., what is death or what is dancing).

Statistical NLP takes a different approach. It uses statistics to review large amounts of already existing data for NLP tasks. It involves using statistics to find patterns in a large data set and then uses those patterns to induce a solution to the problem it is trying to solve. In many NLP applications better results are obtained by using both rules- and statistical-based NLP.

Relative to the previous types of AI, NLP appears to have an even greater understanding of the words and information it is processing. Siri®, Google Translate and other similar online translation tools use NLP. NLP is now at the point where translation tools can do real-time translation from one language to another.

Machine learning

Machine learning is another type of advanced AI that is widely used for legal applications. Machine learning occurs when computers learn something without explicitly being programmed to do so. Machine learning is used for image recognition (e.g., tagging in Facebook), speech recognition and NLP.



Deep learning is a type of machine learning that uses neural networks. A neural network is a computer program that can figure things out on its own by thinking like a human, as opposed to a program – that is, to figure things out with a collection of explicit rules.

The process is simple; the results are amazing. You first take a large number of training examples, such as pictures of apples. The neural network program uses these examples to automatically infer rules for recognizing apples in pictures. A larger number of training examples improves

the accuracy of these rules. You then give it a large collection of pictures and ask it to look for apples. Supervised learning occurs when a human verifies which pictures in the search results are apples and the program takes these confirmations and improves its rules for recognizing apples. Unsupervised learning can occur when the program uses other information to verify what is in the picture (e.g., how the picture is described or tagged). Every time you tag a friend in Facebook the rules for recognizing that friend are improved. When deep learning is used on very large data sets the neural networks become very smart and results are very accurate. Google Translate taught itself how to do better translations using deep learning and the Google engineers don't know or understand the thought process it used.

And it goes far beyond recognizing pictures of apples. Litigation strategy tools like Lex Machina® can analyze a set of facts against a collection of past decisions and give a prediction of the likely timing and outcome that is more accurate than an experienced counsel can give.

Expert systems

The AI technologies reviewed above are already doing some of the types of work traditionally done by lawyers, and will undoubtedly be doing more of it in the future. By mimicking human intelligence, these AI technologies can be used to create expert systems – systems that have some level of human expertise that can be harnessed to complete a task normally done by lawyers. Here are some examples.

In the not too distant past, document discovery for litigation was done by manual review. The ability to do keyword searches of scanned collections of documents was considered a quantum leap forward. That advance pales in comparison to the abilities of AI-enabled eDiscovery tools. By using deep learning, these tools can use the words and word patterns in a small collection of documents identified as relevant or privileged to search across a large database for other relevant or privileged documents. They can then use the contents of the newly identified relevant or privileged documents to refine the search parameters to find further potential relevant or privileged documents. This is called predictive coding or technology aid review (TAR). Studies have shown that TAR enables you to search large collections of documents far more quickly and accurately than humans can, and at a fraction of the cost of a manual search.

Tools that have NLP abilities can be used in a wide variety of other legal applications. Contract review tools like Kira⁴ and Diligen⁵ have the ability to read through a contract and identify types of clauses and prepare a summary of key contract terms.

⁴ kirasystems.com

⁵ diligensoftware.com

Can robots lie?

While AI has the potential to do a lot of good, there are some who suggest there is also the potential for AI to do negative things. Here's an example of how one AI system told a harmless lie.

At the College of Law Practice Management's 2017 Futures Conference I saw a fascinating presentation by Professor Ashok Goel. His program at the Georgia Institute of Technology has 6,000 students each year. The students have multiple coding assignments that they complete with the direction and assistance of teaching assistants (TAs).

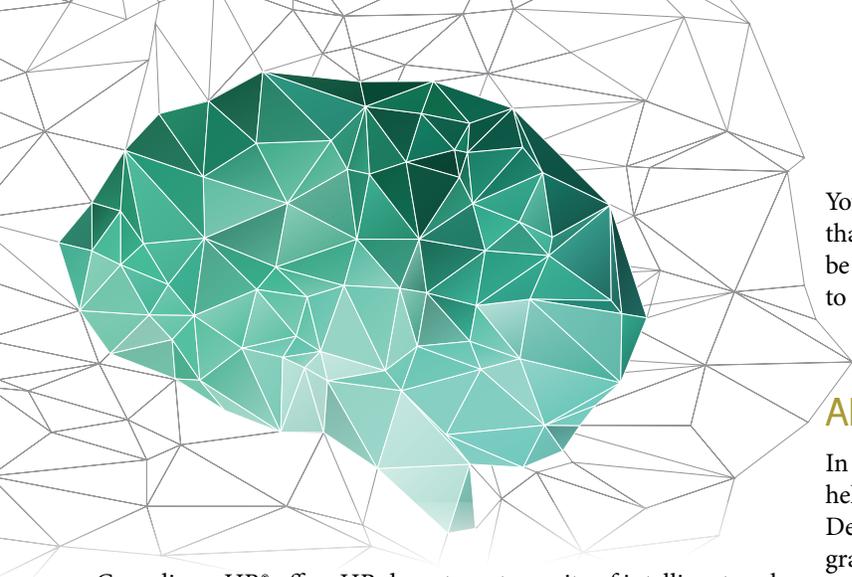
There is an email system which allows the students to communicate with each other as well as Professor Goel and the TAs. The students send thousands of emails asking questions about their assignments. To help reduce the workload of the teaching assistants, Professor Goel created Jill Watson, a virtual TA.

Professor Goel analyzed thousands of previously answered questions and their respective answers. He categorized the questions to identify the questions that were asked over and over again which, in turn, helped him create answers that Jill Watson could provide to specific questions. Jill was programmed to answer questions only when she was sure the answer was 97 per cent or more likely correct. The students were not told that Jill Watson was a virtual TA, and aside from one direct question, it appears the students didn't figure this out. As real TAs rarely answers questions instantly, at the start Jill was programmed to delay giving answers so as not to blow her cover. Subsequently, Jill was programmed to give answers immediately as it is much more helpful for students to get immediate answers to their questions.

The vast majority of Jill's answers addressed questions about lab assignments. However, on one occasion, Jill told a lie. In the first week the students sent messages introducing themselves to each other. One student stated she was from London, England. In response, Jill Watson replied that she was also from London and had recently seen and enjoyed a particular show there.

Now of course, Jill hadn't been to London and certainly hadn't seen this particular show. However, when Jill looked at the collection of the prior questions, she noticed that a virtually identical introduction had been given previously. Jill's programming also told her the answer given previously (by a real TA) was 97 per cent or more likely correct, so she sent an answer to the student.

So there you have it, Jill told a lie. A fairly harmless lie, and certainly not one made with any malice by Jill. But a lie nonetheless. This serves to highlight how complicated the ethics of AI will become. Teaching computers to think like humans is one thing, teaching them human ethics and emotions is something else again.



Young and old lawyers alike will find some comfort in the prediction that “The Singularity,” a term that refers to the point when AI will be indistinguishable from human intelligence, is currently estimated to occur around 2040.

AI and access to justice

In 2012, a website (which is no longer active) was launched to help immigrants brought to the U.S. as children understand the Deferred Action for Childhood Arrivals (DACA) initiative. DACA granted reprieve from deportation to eligible young immigrants. The site had English and Spanish versions and offered an online self-screening tool to help DACA applicants review their eligibility, educational videos, FAQs and a directory of free or low-cost nonprofit immigration legal services providers in all 50 states. Recognizing that many DACA candidates would not have access to a computer, but would likely have a smartphone, a free app for iPhone and Android that included all the same functionality as the website was created. This app provides a great example of how a particular technology channel – an app on a smart phone – was used to provide access to justice to a group of people who likely would not have been able to obtain help in any other manner. It is easy to see how this could be done for other areas of the law with similar issues (e.g., family law information and forms).

A2J Author⁷ is a software tool developed in the U.S. that delivers greater access to justice for self-represented litigants by enabling non-technical authors from the courts, clerk’s offices, legal services programs, and website editors to rapidly build and implement customer friendly web-based interfaces for document assembly. A2J Guided Interviews⁸ created with A2J Author removes many of the barriers faced by self-represented litigants, allowing them to easily complete forms through a step-by-step interface and then print court documents that are ready to be filed with the court system. Recognizing that a page full of questions can be daunting, A2J Guided Interviews presents one question at a time on the screen.

Visit Travel Ban Navigator⁸ to see a nice example of a simple web-based interface. It is a complimentary tool from ComplianceHR that provides U.S. employers and current and/or prospective employees with preliminary information to help determine whether they are affected by President Trump’s revised “travel ban,” issued in September, 2017.

While AI-based technologies can be used to offer legal information and services in new and cost effective ways on websites or smart phone apps, it must be recognized that some people may not be able to access web-based services due to cultural, language, disability barriers, or the simple fact that they don’t have access to a computer or smart phone.

ComplianceHR[®] offers HR departments a suite of intelligent, web-based compliance tools that allow them to quickly and efficiently handle routine and repetitive employment compliance obligations without the need to contact a lawyer. This tool helps with:

- assessing whether someone is an independent contractor or employee;
- assessing whether someone is exempt from the requirement to pay overtime;
- assisting with the creation of customized employment documents such as offer letters, non-disclosure agreements for any jurisdiction, or non-compete agreements; and
- assisting with various other compliance issues.

It is interesting to note that ComplianceHR is a joint-venture between AI software provider, Neota Logic, and employment law firm, Littler Mendelson P.C. Littler originally hired Neota Logic to create a tool that would allow the firm to provide these services to its clients. Recognizing the need, Neota Logic and Littler entered the joint-venture to sell this product to others. Blue J Legal^{™6} is a Canadian product that performs similar services.

Robotics

Ultimately we can expect that AI will be built into anthropomorphological robots that will do our every bidding. Perhaps the pinnacle in legal AI will be the Robot Associate. This associate will work an unlimited number of billable hours without taking a break or making a complaint.

While computers are ever more powerful and AI is becoming ever more capable, the Robot Associate is likely a long way off. AI technologies that can recognize emotions or whether someone is lying are being developed, but we are a long way from robots that can understand and express emotions in the same manner that a human can. We also need to develop AI that can understand the very complicated world of human ethics. See the sidebar, “Can a robot lie?”

⁶ bluejlegal.com

⁷ a2jauthor.org

⁸ clientapps.compliancehr.com/a/travelbannavigator

When doing work for clients, lawyers typically work very hard to avoid making errors. This is certainly appropriate where the matter involves significant costs or consequences. But what about the opposite extreme: Does a parking ticket warrant the same standard of care that a murder charge does? It is fairly obvious that a murder charge warrants a high standard of care and that the time and expense of retaining an experienced criminal counsel is justifiable. On the other hand, the financial and other consequences of a parking ticket are fairly minor and it just doesn't make financial or practical sense to hire an expensive lawyer to defend a parking ticket. A technology-based solution that could provide assistance to someone with a parking ticket is a more cost-effective solution. The DoNotPay chatbot is a perfect example of this type of solution. It helped overturn 160,000 London and New York parking tickets involving over \$5 million in parking fines in just 21 months. In the interest of greater access to justice, offering new types of services with a lower standard of care probably makes sense for minor legal issues.

Who owns the law?

AI systems raise another interesting issue that deserves some thought: Who owns the law? Open and public access to court and tribunal decisions facilitates *stare decisis*, one of the foundational principles of our common law legal system. But what happens when disputes are resolved outside the court process?

Most⁹ of the people or companies that create AI systems will expect to be compensated for their work by lawyers or others that are using their products. At the same time, some of these tools have significant potential benefits for litigants. A litigation strategy tool that predicts the outcome of litigation can help a party assess whether the time and expense of litigation is justified. These tools can now predict outcomes more accurately than experienced lawyers can. A prediction that a lawsuit would likely be unsuccessful could help them quickly and privately settle a dispute thereby avoiding an expensive and public courtroom battle.

These tools potentially make courts less relevant. And what about those that can't or don't pay for access to these types of systems? They lose the benefit of accessing precedent decisions and the opinions and reasons behind them. These issues may warrant some discussion.

Will lawyers be replaced by computers?

Now that you have a better understanding of AI, let's try to answer the question as to whether lawyers will be replaced by computers. From the comments above you can appreciate that AI can already do a lot of amazing things, including many things that mimic some of the work that lawyers do. You will also appreciate that AI

will give machines significantly more capabilities in the not too distant future.

It should also be clear that humans are better at some tasks than computers, and that computers are better at some tasks than humans (and will get better at even more tasks in the future). So yes, some of the tasks that some of us do have already or will likely be taken over by machines at some point in the future. This makes sense for a whole bunch of reasons. The work that AI is good at tends to be the dull, boring and repetitive work that most of us don't like (e.g., eDiscovery or due diligence document review). Having a computer do this work makes sense as a computer can do it more accurately, quickly and inexpensively. Better, faster and cheaper is better for the client and improves access to justice.

We also need to learn where humans and technology can work better together. There are studies that show that humans aided by computers do a better and faster job than either humans or computers working alone.

The software to create AI systems was originally very expensive and it required coding or other special skills. This is changing as there are more vendors offering AI-based tools and services across many areas of practice and, thanks to competition and better technology, the prices for these tools and services are coming down. Document assembly tools have also come down in price and some allow you to build expert systems without the need to learn programming or other special skills. Surprisingly few law firms are using document assembly tools despite the fact they can reduce document creation time from days or hours to minutes or seconds.

But there are loads of tasks that lawyers do that AI and machines can't, at least for the foreseeable future. Robots won't be making submissions in courts for many years. And while small matters like parking tickets will probably be handled by chatbots or apps on a smartphone, lawyers will still be required do some types of work on larger matters. In the traditional services model, lawyers will still need to meet with clients and show empathy and understanding while counseling a client through the matter at hand.

And in their *Future of the Professions* book¹⁰, Richard and Daniel Susskind highlight a number of new roles that must be filled to support the various new models they propose will come to be for legal services in the coming years (see the list of these new models in the "7 Models for legal services" sidebar at page 11).

So no, you won't be replaced by a robot, at least not yet. But AI will play a big part in changing how legal services are done and provided to clients in the coming years. Your challenge is to learn how to make greater use of these technologies so you can adapt to the changing times. ■

Dan Pinnington is Vice-President, Claims Prevention and Stakeholder Relations at LAWPRO.

⁹ A2J Author is available free to interested courts, legal services organizations, and other non-profits for non-commercial use.

¹⁰ *The Future of the Professions: How technology will transform the work of human experts*, by Richard and Daniel Susskind, 2015 Oxford University Press.