Artificial Intelligence and its Application to Law and Legal Services

By Dera J. Nevin[1]

In the popular imagination, artificial intelligence (AI) evokes the machines of the *Matrix* trilogy, the lovable robot depicted in Wall-E or *2001's* Hal. However, depictions of AI are not a recent imaginative invention. Early artificial thinking beings appear in Greek myths, such as Pygmalion's Galatea, and fiction, such as Mary Shelley's *Frankenstein*. Our representations of AI often reveal our cultural conflict between the promise of technological innovation, and the fear of replacement – or enslavement! – by machines. AI also raises a host of ethical and moral issues, which, in addition to the popular conceptions of AI influencing perception, and that AI ranges across so many things, complicates discussion of the topic.

What is AI?

AI is the intelligence exhibited by machines or software. AI is also the name of the field of the study and design of intelligent agents, which are systems that perceive their environment and take actions towards a defined goal. The field has, as a general proposition, that human intelligence can be so precisely described that a machine can be made to simulate it.[2] The AI field is interdisciplinary, and encompasses computer science, mathematics, neuroscience, linguistics, philosophy and psychology. There are multiple branches of AI, with various disciplines pursuing different capabilities as varied as: understanding human speech, interpreting images and objects, competing in games, interpreting complex data, and driving a car. Because there are so many objectives, and different scientific, mathematical and computing approaches to achieving those objectives, the study of AI and its commercialization efforts are divided into subfields that focus on specific problems or applications. I have found it helpful to think of AI as a cluster of technologies, approaches and disciplines.

Modern AI research aims to reproduce "human" cognitive behavior using machines, including the sub-problems of: knowledge, planning, learning, reasoning, communication, perception and interaction with objects (and then humans). The solution of one or more of these problem subsets is generally described a narrow (or weak) AI.

Most of us experience narrow (weak) AI on a daily basis, simply because it is already all around us, most of the time. John McCarthy, who is credited with coining the term "artificial intelligence" in 1956, noted that "as soon as it works, no one calls it AI anymore."[3] Examples of narrow AI include: the spam filter in email, the recommendation engines in Amazon and Netflix, Siri, most GPS navigation technology and the Kinect system for the Xbox 360 and the Xbox One. Other examples of narrow AI systems are IBM's Deep Blue, a computer chess-playing computer system which in 1997 defeated Garry Kasparov, and IBM's

---

[1] eDiscovery Counsel, Proskauer Rose LLP. The opinions expressed in this article are her own. She would like to thank Ian Kerr and Noah Waisberg for the multiple opportunities to discuss these issues and AI technology and its application to the delivery of legal services. Elements of the information in this paper have previously been published by the Canadian Bar Association and have been delivered in plenary sessions for the Canadian Bar Association (August 15, 2015), and the International Legal Technology Association (September 1, 2016).

[2] I owe this understanding of the conjecture underlying AI research to Stuart J Russell and Peter Norvig, *Artificial Intelligence, A Modern Approach (*2nd Ed.) (2003), Prentice Hall, and their discussion of "The Dartmouth proposal".

[3] Unattributed quote, cited by Moshe Y. Vardi, "Artificial Intelligence: Past and Future," *Communications of the ACM, Vol. 55 No. 1, Page 5, describing the 1956* Dartmouth College conference in which the field of AI was "founded".

Watson, a computer that can answer questions using a sophisticated natural language processor, which in 2011 defeated *Jeopardy*! champion Ken Jennings. More recently, Google's AlphaGO,[4] and its successor, Master,[5] have defeated national and international champions at Go, which is a game that many consider to be more complicated than chess. Google Translate and self-driving cars are further examples of ANI systems (often multiple ANI systems working in conjunction).

AI in the Legal Context

Two of the most talked-about uses of narrow (weak) AI in the legal field involve machine learning and natural language processing.

Machine learning involves the development of computer algorithms that improve 'automatically' through experience and repetition. Using unsupervised learning, a computer finds patterns in data without human input or intervention. Using supervised learning, classification is used to determine what category data belongs to after a computer is presented with a number of examples by a human in training rounds; the computer then uses regression (a mathematical computation) to produce a function (calculation) that accounts for the relationship between the inputs and outputs and predicts how the outputs could change as the inputs change. Machine learning applications are becoming prevalent in eDiscovery technology, variously described as "predictive coding", or "technology-assisted review".

Natural language processing (NLP) involves computers reading and understanding human language, and a sufficiently developed NLP interface enables a machine to obtain 'knowledge' (information) directly from written and spoken sources. Existing NLP applications include text mining, machine translation and answering questions. Among the first applications of NLP In law was legal research, with offerings from Westlaw and Lexis about a decade ago. More recently, Fastcase, among others, has added visualization features to the information delivered by NLP, since pictures and visual relationships can help humans better understand the results.

IBM's *Watson* is an example of a question answering machine with a powerful NLP interface. ROSS (ROSS Intelligence) is the application of IBM's Watson to legal research (more specifically, as a first offering, bankruptcy information). When ROSS is pointed at a body of legal texts and statutes, and presented with a question, its programming permits it to "answer" the question, by using NLP and other weak AI components to perform the legal research and present information back in a process similar to a recommendation engine Amazon or Netflix); as the human interacts with the information delivered by ROSS, this serves to improve the subsequent recommendations. Other companies have announced collaborations with Watson and we could soon see additional commercial applications of this technology to large sources of caselaw and other data about law.[6]

---

[4] Damis Hassabis, "AlphaGo: using machine learning to master the ancient game of Go", https://blog.google/topics/machine-learning/alphago-machine-learning-game-go/ (1/11/2017)

[5] "Google's AlphaGo AI secretively won more than 50 straight games against the world's top Go players", Quartz, https://qz.com/877721/the-ai-master-bested-the-worlds-top-go-players-and-then-revealed-itself-as-googles-alphago-in-disguise/ (1/11/2017)

[6] The technology of IBM *Watson* works best with a big data (i.e. a big source of mostly-organized data, such as a repository of caselaw, or statutes), and it also requires "training" so that the system has sufficient inputs to make recommendations. That legal commercial applications are taking time to evolve reflects that it takes human expertise and a lot of time and effort to assemble the necessary data sets required, analyze the content and the information retrieved by the system, and test the results to ensure the system operates optimally.

Similarly, machine learning and NLP have been combined in Contract Analysis applications such as Kira Systems and KM Standards. The contract portfolio of an organization can be "fed" into one of these systems which, based on the previous training of these systems, can identify standard clause language and contract structure, improving due diligence outcomes and simplifying compliance review. "ACE", by RAVN, will "read, interpret and summarize" key information from contracts. A related technology, Neota Logic, combines an expert system with NLP and machine learning techniques to accomplish tasks as varied as input-dependent workflow, and question-answering.

Finally, technologies underlying legal analytics (outcome prediction) were highlighted when Lexis Nexis acquired Lex Machina in November 2015. Lex Machina assembled a set of IP cases and decisions which it mined using NLP and other predictive analytic techniques to predict outcomes of IP litigation. LexPredict points its predictive models to the Supreme Court jurisprudence. Other offerings are pointing directly to court dockets to expose decision patterns in judges and lawyers' win/lose histories.

Machine learning and NLP such as ROSS using IBM Watson will be powerful game changers in legal. In a 2011 study titled "Where the Money Goes", the RAND Corporation studied eDiscovery outcomes and noted the use of predictive coding would substantially lower eDiscovery costs.[7] Similarly, ROSS has been described as game-changing for legal research, particularly if commercialized for individual litigants, although due to the effort involved in training the system, it may be several years before we see this. In eDiscovery, document review roles are being eliminated and some observe that certain lawyers' work in research may similarly disappear if ROSS becomes prevalent.

More likely, AI technology will reduce the time lawyers currently spend reviewing documents of low evidentiary value and eliminate duplicative legal research, freeing up lawyers to do other work. However, at its current stage of evolution, AI generally does not replace the combination of legal skills required to provide legal advice or resolve many legal issues.

General AI is the Goal of AI research

General (or strong) AI remains the field's long-term goal. We do not yet have general or strong AI (at least not in a commercially accessible form) and the academic community is deeply divided on when, if ever, we will. General intelligence would combine all the skills listed above, and may exceed human abilities at most of all of them. General intelligence may require anthropomorphic features like artificial consciousness or an artificial brain. Achieving strong AI is hard, based on the current understanding of how the human brain works, and available computing power. It is actually easier to build a machine that can compute complex equations than it is to build one that can consistently recognize a picture of a dog, let alone consistently recognize different pictures of different dogs across breeds. Currently, a machine would beat me at chess, but could not understand the meaning of the words in this paragraph.

Consider machine translation: the machine must read and write in both languages (natural language processing), know what the text is saying (knowledge), understand the author's argument (reason) and reproduce the text's intention. Anyone who has used current machine translation technologies knows there is room for improvement.

---

[7] "Where the Money Goes", RAND Corporation, http://www.rand.org/content/dam/rand/pubs/monographs/2012/RAND_MG1208.pdf (1/11/2017)

Impact of AI for Lawyers

Will machines exhibiting either weak or strong AI completely replace lawyers? I do not think so, at the current state of technological evolution, although certain tasks currently performed by humans in the course of legal services delivery may be replaced by automation, including AI, and in some contexts reasonably quickly. I can make two predictions.

First, and more generally, AI technologies will change what we do, because it will raise a number of novel legal problems and disputes that lawyers must become equipped to handle. Are you ready to develop policies to protect our privacy from text -mining applications, litigate a case on liability associated with a self-driving car or other automated robot, or draft a contract that is self-executing (or example, one that uses Blockchain)? These days are already upon us.

Second, it may change how to work, as we may have to adapt and use different tools and technology to do our work. This adaptation may be harder: new technology can be challenging, even as it is beneficial, because change, even if welcome, can be experienced as difficult and stressful, and change management is hard, even when change is planned. Many lawyers do not routinely follow technology news, and so experience technology innovation as unplanned, and therefore unwelcome, or become attached to solutions to past problems that have been mostly rendered moot by the new technologies.

Almost always, new technology itself is neutral, and so much depends on its application. Technology generally serves to enlarge us, to extend our reach outside ourselves. Fire helped us warm ourselves when the sun was not present. Boats and the wheel extended our range and sphere of motion. The lens enabled us to see far and see close. The telephone helps us hear far away, to speak to someone not present. These technologies improve the range and strength of our physical bodies and of our senses. Other technologies ingrate with our mental and imaginative faculties; many directly extend the reach and power of our brains. Early computers extended computational speed and accuracy. Portable devices extend (though perhaps do not improve) our memories by externalization of information. Artificial intelligence applications may extend our ability to access repositories of information. All of this will happen whether we as lawyers wish this or not, because it is happening in society.

I won't pretend to know what the future will look like, but I do know the future needs us as lawyers to embrace and incorporate technology into the law and our practice of it. I have understood law to be a technology for organizing society, and a technology that has been proven to be adaptive and flexible. I have observed that societies we describe as lawless often lack basic technology. I have developed the philosophy in my own practice, that to embrace new technologies in the delivery of legal services is to improve the very technology I chose as my calling.

Dera J. Nevin is eDiscovery Counsel at Proskauer Rose LLP. Elements and ideas within this article were first published in The National *magazine (September 2015). The opinions expressed in this article are her own. She can be contacted at dnevin@proskauer.com.*