From Countess Lovelace to Ross

A brief overview of artificial intelligence (and its increasing use in the legal profession)

By Farid Assaf SC – Banco Chambers[1]

Introduction

In the annals of the history of computing, a more unlikely collaboration would be difficult to imagine. On the one hand Charles Babbage (1791 – 1871), the English polymath widely acknowledged as the father of modern computers and on the other Ada, Countess of Lovelace, the (only legitimate) daughter of famed English poet Lord Byron and considered by many to be the first ever computer programmer. In 1833 Babbage demonstrated his newly constructed Difference Engine (a mechanical calculator) to Lovelace and her mother (whom Byron dubbed the 'Princess of Parallelograms'). Lovelace was inspired by what she saw and went on to become an esteemed mathematician in her own right.

Her most famous contribution was the translation of a paper from French to English written by Luigi Menabrea in 1842 (who besides becoming prime minister of Italy also examined the mathematics of structural analysis). In the paper, Menabrea discussed Babbage's 'Analytical Engine' – a successor to the Difference Engine and considered to be the first ever computer (at least conceptually). Not content with simply translating the paper, Lovelace prepared detailed notes of her own which included a suggested algorithm that could program the engine.[1]

Lovelace's notes also contained an observation which some have considered to be a dismissal of artificial intelligence. She wrote: 'The Analytical Engine has no pretensions whatever to *originate* anything. It can do *whatever we know how to order it* to perform. It can follow analysis; but it has no power of anticipating any analytical relations or truths.' In his seminal paper Computing Machinery and Intelligence, Alan Turing sought to countenance Lovelace's perceived negativity towards artificial intelligence (AI). In that paper, Turing posed the question 'Can machines think?' To answer the question, Turing devised a test for artificial intelligence whereby a machine attempts to convince a human interrogator it really is

human through a series of written responses to various questions.

The so-called Turing Test has been criticised by researchers in artificial intelligence. Instead, some researchers suggest the Lovelace Test (named of course after the Countess). An artificial agent, designed by a human, passes the Lovelace Test only if it originates a program that it was not engineered to produce. In other words, the Lovelace Test requires a computer to create something original, all by itself.

There has been a significant increase in AI research since the publication of Turing's paper in 1950. In its inaugural 2017 AI Index, Stanford University estimates that the number of AI research papers produced each year since 1996 has increased more than nine-fold, AI class enrollment at Stanford during the same time-frame has increased eleven-fold and there are now fourteen times the number of active US startups developing AI systems than there were in 2000.[2] Notwithstanding this increase in activity, AI researchers are yet to develop technology which passes the Lovelace Test (or for that matter the Turing Test, at least on a consistent basis). Even so, the surge in interest in AI has pervaded a significant number of industries including the law. AI technologies are now routinely used in a wide variety of industries including health, finance and teaching.

This article briefly explores the current status of AI, its future development and its possible practical uses for the legal profession.

What is AI?

Despite the term 'artificial intelligence' being coined in 1956 by American computer scientist John McCarthy, there is no universal definition of the expression. At the risk of over-simplification, AI can simply be described as 'non-biological intelligence'.[3] The Oxford dictionary defines AI as 'the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.' This definition refers to the various metrics used by researchers to determine whether a non-biological entity truly exhibits artificial intelligence. As will be explained below researchers have been able to build technologies that display some of these attributes but are far from achieving general human level intelligence.

The state of AI technology and predicted future advancements

For present purposes, three broad phases can be described for AI and its evolution: now, near and next.[4] The present epoch is described as one of narrow AI. Narrow AI technologies focus on a limited task designed to replicate and surpass human intelligence.[5] Currently, AI systems already outperform human intelligence in many domains[6] including defeating human champions in a wide variety of games such as checkers (1994 with the CHINOOK program); backgammon (1979 with the BKG program); chess (1997, Big Blue against Garry Kasparov); scrabble and more recently Jeopardy! in 2011. Jeopardy! is a television game show with trivia questions covering a variety of topics including history, geography and literature. In 2011, IBM's Watson defeated two former winners of the game-show.[7]

The near phase encompasses artificial general intelligence (AGI) or human level machine intelligence (HLMI) which is yet to be developed and is defined as AI 'capable of performing all intellectual tasks that a

human brain can.'[8] Estimates as to when this next phase will commence vary, however approximately 50 per cent of recently surveyed AI researchers predict the year 2040 as when that milestone is likely to be reached, whereas 90 per cent of researchers predict AGI would be reached by about 2075.[9] The next phase is artificial superintelligence (ASI). In his *New York Times* best-selling book *Superintelligence*, Nick Bostrom, a professor of philosophy at Oxford University, predicts that superintelligence will be achieved 'relatively soon after' achieving HLMI.[10] In the book he opines that a 'plausible default outcome' of the creation of machine superintelligence is 'existential catastrophe.'[11] The future is not entirely dystopian with Bostrom suggesting humanity could utilise what he calls 'indirect normativity' to effectively delegate to a superintelligence the reasoning required to select certain universal (benevolent) values. He summarises such an approach as a heuristic principle which he labels 'epistemic deference', that is, a superintelligence which 'occupies an epistemically superior vantage point: its beliefs … are more likely than ours to be true. We should therefore defer to the superintelligence's opinion whenever feasible.'[12]

Artificial or 'augmented' intelligence?

During the present transitional phase to HLMI, narrow AI technologies are producing what some have labelled augmented intelligence (sometimes referred to as intelligence amplification). The goal of augmented intelligence is not to replace humans, but rather capitalise on the combination of algorithms, machine-learning and data science to inform human decision-making abilities.[13] IBM is at the forefront of this augmented intelligence research and development. Rob High, Vice President and CTO of IBM Watson, explains IBM's approach in the following way: 'If you look at almost every other tool that has ever been created, our tools tend to be most valuable when they're amplifying us, when they're extending our reach, when they're increasing our strength, when they're allowing us to do things that we can't do by ourselves as human beings. That's really the way that we need to be thinking about AI as well, and to the extent that we actually call it augmented intelligence, not artificial intelligence.'[14] One of IBM's key technologies in developing augmented intelligence is its Watson technology. IBM describes Watson as a cognitive system built on the current era of programmatic computing which utilises deep natural language processing. The uniqueness of the Watson technology is to combine the capabilities of natural language processing (by helping to understand the complexities of unstructured data); hypothesis generation and evaluation (by applying advanced analytics to weigh and evaluate a panel of responses based on only relevant evidence) and dynamic learning (by helping to improve learning based on outcomes to get smarter with each iteration and interaction).[15] The Watson technology is currently used in a wide array of industries including medicine, finance and now law.

AI and the legal profession

Academics and entrepreneurs have identified a number of AI technologies suitable for use in the law and legal practice. In his 2017 book *Artificial Intelligence and Legal Analytics* published by Cambridge University Press, Kevin D. Ashley, a Professor of Law and Intelligent Systems at the University of Pittsburgh, explains and explores the AI systems currently available and that can be specifically adopted for legal work. Ashley explains that the goal of much of the research in AI and law has been to develop 'computational models of legal reasoning' (CMLRs) that can make legal arguments and use them to predict outcomes of legal disputes.[16] A subset of CMLRs known as computational models of legal argument (CMLAs) implements a process of legal argumentation as part of their reasoning.[17] While researchers have made significant progress in developing such models some obstacles have arisen. So far for example, the substantive legal knowledge employed by computational models has had to be manually obtained by legal professionals from legal sources. This inability to automatically connect CMLRs directly to legal texts has limited the researchers' ability to apply their programs in real-world legal information retrieval, prediction and decision-making.[18]

However, recent developments in computerised question answering (such as Watson), information extraction from text (which summarises the essential details particular to a given document) and argument mining (which involves automatically identifying argumentative structures within document texts) promise to change that. All three technologies usually rely, at least in part, on applying machine learning to assist programs in processing semantic information in the texts.[19] Another technique which may assist researchers' ability in producing CMLRs for real-world applications is text analytics or text mining. This technique refers to a set of linguistic, statistical and machine learning techniques that model and structure the information content of textual sources for business intelligence, exploratory data analysis, research or investigation.[20] In the legal context, this technique can be applied (which Ashley refers to as 'legal analytics') so as to derive substantively meaningful insights from legal data.[21] Ashley predicts that some CMLRs and CMLAs may soon be linked with text analysis tools to enable the construction of a new generation of legal applications. As Ashley explains, 'CMLRs and CMLAs developed in the AI and law field will employ information extracted automatically from legal texts such as case decisions and statutes to assist humans in answering legal questions, predicting case outcomes, providing explanations, and making arguments for and against legal conclusions."[22] The above concepts are best explained by way of examples.

Examples of AI technology used in law

Predictive Coding

Predictive coding, also known as Technology or Computer Assisted Review (TAR), is a discovery-specific, dialogic application of machine learning technology in which a program develops, applies and refines a predictive document-search model based on search terms, document categorisations and feedback given by human case managers.[23] TAR uses machine learning to identify relevant documents. The process involves a small team initially reviewing a seed set of documents. Once complete, a computer identifies similarities and patterns within the entire set of documents and attempts

to determine coding that will be useful for additional document sets.[24] The United States has been an early adopter of TAR. In *Da Silva Moore v Publicis Groupe et al* (2012) 287 F.R.D 182, Judge Peck, a then federal magistrate judge for the United States District Court for the Southern District of New York, observed, in an opinion specifically addressing the issue, that TAR was the best methodology to process the nearly three million-odd documents in that case (a sex discrimination case) as opposed to manual review. Numerous other US cases have taken a similar approach as have English courts (see for example the 2016 case of *Pyrrho Investments Limited & Anr v MWB Property Limited* [25]). TAR has also been adopted in Australia. In *McConnell Dowell Constructions (Aust) Pty Ltd v Santam Ltd (No 1)* [2016] VSC 734 the plaintiff identified approximately four million potentially relevant documents. After referring to the US and English cases mentioned above, Vickery J of the Supreme Court of Victoria endorsed the suggestion of a court-appointed special referee that the parties use TAR to increase the efficiency of the document review. His Honour noted that the use of technology in civil litigation should facilitate the just, efficient, timely and cost-effective resolution of the real issues in dispute as required by section 9 of the *Civil Procedure Act 2010* (Vic).[26] The Supreme Court of Victoria has also issued a specific practice note dedicated to the use of technology in civil litigation[27] where the court makes clear that practitioners are expected to consider the use of technology as early as possible. The Federal Court has also issued a similar practice note.[28]

Ravel Law and Lex Machina

Ravel Law is a startup by two former Stanford law students who have sought to disrupt traditional text intensive legal research. One of the most distinctive features of Ravel's user interface is the display of legal research search results – rather than appearing as blocks of text, search results appear as an interactive visualisation. Case results are displayed as bubbles of various sizes – landmark cases are depicted as larger bubbles while less important cases appear smaller. The relationship between the various case bubbles are depicted graphically. Search results can also be filtered in a number of ways such as rulings from various courts or dates. Another distinctive aspect of Ravel's platform is its analytics suite which includes court analytics; judge analytics; and case analytics.[29] The analytics suite analyses case data to produce practical summaries. For example, using the judge analytics function users can see how a particular judge will respond to a particular application based on past data involving relevant factors.[30] For example: 'Judge Susan Illston in the Northern District of California grants 60 per cent of motions to dismiss, which makes her 14 per cent more likely to grant than other judges in the district.' At the time of writing the Ravel Law technology is available to LexisNexis Advance users in the United States.

Similar to Ravel Law, Lex Machina uses natural language processing to analyse court documents that are publicly available to try to predict things like the ruling of a particular judge in a particular case or the behaviour of a particular lawyer.[31] Lex Machina uses natural language processing to analyse court documents that are publicly available to try to predict matters such as the ruling of a particular judge in a particular case, the behaviour of a particular lawyer and the litigation history of particular parties.[32] Lex Machina originated with a particular focus on IP, however, since its acquisition by LexisNexis, Lex Machina has now branched out into other practice areas such as trademark and copyright litigation.[33]

"If to err is human, how do I explain this mess."

United Archives GmbH / Alamy Stock Photo

Ross

Based on question-answering computer system IBM Watson, ROSS is a cloud-based system that uses natural language processing and machine learning capabilities to understand, research and provide answers to legal research questions.[34] IBM's Watson, upon which ROSS is based, is in turn reliant on what is referred to as Unstructured Information Management Architecture ('UIMA'), which is a framework which uses a series of software components called annotators to analyse text and draw increasingly abstract inferences about textual meaning.[35] Upon provision of an answer, the human interacting with ROSS then tells the system whether the answer provided was relevant and ROSS uses this information to learn to produce just as relevant or more relevant answers in the future.[36]ROSS's current capability extends to bankruptcy, intellectual property and labor and employment law in the United States.

IBM's Debater

IBM describes Project Debater as, 'the first AI system that can debate humans on complex topics.' IBM says that Project Debater relies on three 'pioneering capabilities.' First, data-driven speech writing and delivery, which is said to be the ability to automatically generate a whole speech and deliver it persuasively. Secondly, is listening comprehension, which IBM describes as the ability to understand a long spontaneous speech made by the human opponent in order to construct a meaningful rebuttal. Thirdly, is the system's ability to model human dilemmas and form principled arguments made by humans in different debates based on a unique knowledge graph. IBM claims that by combining these core capabilities it can conduct a meaningful debate with human debaters. The development of Project Debater has required IBM to venture into new and discrete areas of AI research such as argument mining (i.e. identifying an argument and its position with respect to the relevant topic); debate speech analysis (which entails the ability to understand and rebut the text of the opponent's speech and the development of text to speech systems (i.e. the ability to interact with its surroundings in a human-like manner). For readers who are interested, IBM has released datasets for Project Debater which sets out the various comprehensive research papers relied upon for the project. The link to those datasets is contained in the end notes to this article.[37]

In a presentation held in San Francisco on 18 June 2018, IBM demonstrated the Debater technology with a real-time debate with a human on the topic of whether government should subsidise space exploration. IBM's

debater had no awareness of the debate topic ahead of time. Each side gave a four-minute introductory speech, a four-minute rebuttal to the other's arguments and a two-minute closing statement. According to media reports, the AI Debater 'held its own.'[38] During the debate, the AI cited sources, indulged the audience's affinity for children and war veterans, utilised analogies and even made a few 'passable' jokes. Having viewed extracts of the debate[39], the writer can confidently say that members of the New South Wales Bar should not be concerned about job security – at least at this stage.

Data Privacy Advisor - Thomson Reuters

Data Privacy Advisor is another AI tool based upon Watson technology. Launched earlier this year, the Advisor is a tool primarily designed to assist compliance officers keep up to date with the myriad of privacy regulations faced by businesses around the world. The tool contains global statutory and regulatory data privacy country guides for more than 80 countries, question answering capability using natural language through IBM Watson-enabled technology as well as curated news, analysis and blog content specific to data privacy.

Conclusion

It is hoped that this article has provided some insight into the world of AI and its application in the legal profession and that the popular cacophony[40] regarding AI can be placed into context. At this stage of research, and at least for the foreseeable future, the legal profession, and professionals, are likely to be greatly assisted by the various augmented intelligence technologies being developed. We are a long way from technology being able to satisfy Countess Lovelace's quite demanding test for AI. The literature suggests however that it is almost inevitable that day will arrive although it is likely to be towards the end of this century. The consequences of this for the legal profession will need to be examined in a subsequent article.

Gualtiero Boffi / Alamy Stock Photo

ENDNOTES

1 The author thanks Elizabeth Eggerton and Jason Leonard from IBM for their generous time in providing valuable insight into IBM's research on artificial intelligence. Thanks are also extended to Penny Thew of 8 Wentworth for the kind referral to IBM.
2 Artificial Intelligence Index – 2017 Annual Report (http://aiindex.org/2017-report.pdf)
3 Max Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence* (Suffolk: Penguin Books, 2017) at 39.
4 David Skerrett, 'The State of Artificial Intelligence around the Globe' 29 August 2018, http://www.econtentmag.com/Articles/Editorial/Mobile-Moment/The-State-of-Artificial-Intelligence-Around-the-Globe-126115.htm
5 https://searchenterpriseai.techtarget.com/definition/narrow-AI-weak-AI
6 Nick Bostrom. *Superintelligence: Paths, Dangers, Strategies*. Suffolk: Oxford University Press, 2014 at 15 – 17.
7 https://www.nytimes.com/2011/02/17/science/17jeopardy-watson.html
8 David Skerrett, 'The State of Artificial Intelligence around the Globe' 29 August 2018, http://www.econtentmag.com/Articles/Editorial/Mobile-Moment/The-State-of-Artificial-Intelligence-Around-the-Globe-126115.htm
9 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Suffolk: Oxford University Press, 2014 at 23.
10 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Suffolk: Oxford University Press, 2014 at 25.
11 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Suffolk: Oxford University Press, 2014 at 140.
12 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*. Suffolk: Oxford University Press, 2014 at 258.
13 Nick Ismail, *Augment Intelligence: Why the Human Element Can't Be Forgotten* (21 June 2017) Information Age <https://www.information-age.com/augmented-intelligence-human-element-cant-forgotten-123466894/>.
14 Dan Costa, 'IBM Watson CTO on Why Augmented Intelligence Beats AI', 14 August 2017 (https://au.pcmag.com/fast-forward/49251/ibm-watson-cto-on-why-augmented-intelligence-beats-ai)
15 Rob High, 'The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works', 2012 (https://developer.ibm.com/watson/wp-content/uploads/sites/19/2013/11/The-Era-of-Cognitive-Systems-An-Inside-Look-at-IBM-Watson-and-How-it-Works1.pdf)
16 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 4.
17 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 4.
18 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 4.
19 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 4 – 5.
20 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 5.
21 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 5.
22 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 5.
23 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 241.
24 Thomas Davey and Michael Legg, 'Predictive Coding: Machine Learning Disrupts Discovery', (2017) 32 *Law Society of NSW Journal* 82.
25 *Pyrrho Investments Limited & Anr v MWB Property Limited* [2016] EWHC 256 (Ch).
26 Herbert Smith Freehills, 'One Giant Leap for E-Discovery: Predictive Coding approved by an Australian Court' (6 December 2016) <https://www.herbertsmithfreehills.com/latest-thinking/one-giant-leap-for-e-discovery-predictive-coding-approved-by-australian-court>.
27 Supreme Court of Victoria, *Practice Note No SC Gen 5 – Technology in Civil Litigation*, 1 January 2017.
28 See for example Part 4 of the Technology and the Court Practice Note (GPN-TECH).
29 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 354.
30 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 354.
31 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 353; Karen Turner, 'Meet 'ROSS,' the Newly Hired Legal Robot' *Washington Post* (16 May 2016) 1.
32 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 353.
33 Harvard Business School, 'Lex Machina: Moneyball Law' <https://www.hbs.edu/openforum/openforum.hbs.org/goto/challenge/understand-digital-transformation-of-business/lex-machina-moneyball-law.html>.
34 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 351; Andrew Arruda, 'An Ethical Obligation to Use Artificial Intelligence: An Examination of the Use of Artificial Intelligence in Law and the Model Rules of Professional Responsibility' (2017) 40 *American Journal of Trial Advocacy* 443, 451–3.
35 Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Suffolk: Cambridge University Press, 2017 at 202.
36 Andrew Arruda, 'An Ethical Obligation to Use Artificial Intelligence: An Examination of the Use of Artificial Intelligence in Law and the Model Rules of Professional Responsibility' (2017) 40 *American Journal of Trial Advocacy* 443, 453.
37 https://www.research.ibm.com/haifa/dept/vst/debating_data.shtml#Argument%20Detection
38 https://www.theverge.com/2018/6/18/17477686/ibm-project-debater-ai
39 https://www.youtube.com/watch?v=s_wgf75GwCM
40 See Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Suffolk: Oxford University Press, 2014) at 321.