



THE LAW SOCIETY OF NSW  
IN COLLABORATION WITH UNSW LAW

# ARTIFICIAL INTELLIGENCE AND THE LEGAL PROFESSION: A PRIMER

---

Professor Michael Legg  
and Dr Felicity Bell



THE LAW SOCIETY  
OF NEW SOUTH WALES



UNSW  
SYDNEY

## WHAT IS FLIP STREAM?

*A strategic alliance between the Law Society of NSW and UNSW Law aims to tackle the challenges of technological change and its impact on lawyers, law and the legal system.*

In 2016 the Law Society of NSW established the Future Committee and, in turn, the Future of Law and Innovation in the Profession (FLIP) Commission of Inquiry. In March 2017, the inquiry culminated in the Law Society's ground-breaking FLIP Report, which discusses the future of the legal industry in the digital age.

The Report recognised the legal profession is undergoing change at a pace never before experienced and in unforeseen ways. This change has major ramifications for not just the legal profession, but for clients and society more generally, particularly in relation to access to justice.

In November 2017, the Law Society entered into a strategic alliance with University of New South Wales (UNSW) Law to generate a stream of research to consider and respond to the issues raised by the FLIP Report, such as legal technology, clients' needs and expectations, new ways of working, community needs and legal education, artificial intelligence and the practice of law and technological solutions to facilitate improved access to justice.

This dedicated research stream will also tackle some of the increasingly complex challenges presented by digital and other technological transformations and its impact on lawyers, law and the legal system.

This strategic alliance, forged between a world-class university, UNSW, and the Law Society is a milestone of progress for both institutions and for the entire legal profession.

Our organisations are meeting the challenges and opportunities presented by technology and innovation in our operating environment head on, driven by a shared mission:

To help equip Australian lawyers with the tools they need to confront the future with confidence and ease.

Each year the FLIP Stream, as it has become known, will undertake research into an annual topic that will then be disseminated through the academy, the profession and society. In 2018 the annual topic was Artificial Intelligence and the Legal Profession, led by Professor Michael Legg and Dr Felicity Bell. The 2019 topic on Change Management is led by Dr Justine Rogers. The FLIP Stream will also engage in and respond to other areas of research and law reform.

The Law Society is encouraged and excited by this alliance, knowing that our members and the people we serve will be the ultimate benefactors.

# **ARTIFICIAL INTELLIGENCE AND THE LEGAL PROFESSION: A PRIMER**

Professor Michael Legg and Dr Felicity Bell

THE LAW SOCIETY OF NEW SOUTH WALES'S  
FUTURE OF LAW AND INNOVATION IN THE PROFESSION  
RESEARCH STREAM, UNSW LAW (FLIP STREAM)

## INTRODUCTION

In 2017 the Law Society of New South Wales published the findings of its Future of Law and Innovation in the Profession (flip) Commission of Inquiry. A key finding was that legal practices are increasingly interested in and engaging with legal technology. One of those technologies was Artificial Intelligence (AI). However, lawyers' level of understanding and use of technology was uneven across the profession.<sup>1</sup>

AI has existed as a concept since the 1950s and the idea that AI could be applied to the law has been explored since the 1980s. However, over the years, progress in the development of AI has been cyclical, and interest in and funding of AI research has fluctuated, with a number of AI Winters in which progress stagnated. More recently AI has experienced a new Spring with significant leaps forward that have begun to carry over to legal practice.

This primer on Artificial Intelligence and the Legal Profession seeks to introduce the concept of AI and how it may be used in legal practice to lawyers. The discussion is introductory and aimed at raising the level of understanding of AI across the legal profession.

## WHAT IS ARTIFICIAL INTELLIGENCE?

### Definitions

The term artificial intelligence (or AI) may be traced back to the 1950s but the idea of non-humans exhibiting intelligence may be traced back much further in religion and literature. The meaning of artificial intelligence is contested and subject to change. "Artificial" is used to denote something that is not natural, frequently meaning non-human and usually associated with machines: computers or robots. The idea is to indicate where the so-called intelligence comes from, but to distinguish it from human intelligence. "Intelligence" is usually synonymous with reasoning, the ability to learn or understand. However, intelligence and reasoning involve the use of a number of other attributes or skills which when combined are recognised by humans as demonstrating intellect. As explained below AI is not really intelligent in the sense explained above as AI does not know what it is doing, or why it is doing it. An AI system is not really 'reasoning' or 'thinking' but is following a set of pre-programmed computational steps (expert systems) or mathematically analysing a huge amount of data to infer a probability (machine learning). However, the mathematical analysis can be performed autonomously with the AI devising patterns or relationships without human input (unsupervised machine learning).

---

<sup>1</sup> Law Society of New South Wales Commission of Inquiry, flip: The Future of Law and Innovation in the Profession (2017) 5.

AI, as a term or field of computer science, is employed where processes are used to carry out tasks which, if performed by a human, would be seen as evidence of intelligence – i.e. the processes mimic, imitate or simulate intelligence. AI may be defined by reference to the tasks it performs such as visual perception, speech recognition, decision-making, and translation between languages. Alternatively, AI may be defined by reference to the processes used to perform tasks: expert systems, machine learning (supervised, unsupervised, neural networks).<sup>2</sup>

### Strong/Weak and General/Narrow AI

There are different ways of classifying AI: “strong” or “weak”,<sup>3</sup> and “general” or “narrow”, among others. Strong AI refers to AI which can ‘think’ in an independent manner; weak AI refers to a program mimicking human thinking but without actually being able to reason similarly to a human.<sup>4</sup> We can also distinguish between “general” and “narrow” AI.<sup>5</sup> In this primer, we are primarily concerned with narrow AI. AI currently works best with focused, precisely defined tasks, and most current legal applications of AI fall into this category.<sup>6</sup> AI encompasses a number of different branches,<sup>7</sup> including robotics, computer vision and speech functions (see Figure 1). Not all these branches of AI have express application to the law or legal services.

Narrow AI systems may surpass human performance in a specific task but are unable to generalise this capability to other tasks or domains as a human could. IBM’s Deep Blue system and Watson are examples of narrow AI. Neither system would be able to apply its technology to another domain without significant human guidance through reprogramming and data inputs. To enable these systems to perform in a new domain would be ‘the analogue of needing to perform brain surgery on a human each time the person needs to confront a new sort of task’.<sup>8</sup>

---

<sup>2</sup> For discussions of how artificial intelligence may be defined see Viktor Mayer-Schonberger and Kenneth Cukier, *Big Data* (First Mariner Books, 2014); Stuart J Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach* (Pearson, 3<sup>rd</sup> ed. 2016); Jerry Kaplan, *Artificial Intelligence, What Everyone Needs to Know* (Oxford University Press, 2016); Michael Mills, *Artificial Intelligence in Law: The State of Play* (Thomson Reuters, 2016); Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age* (Cambridge University Press, 2017); Jacob Turner, *Robot Rules* (Palgrave Macmillan, 2019).

<sup>3</sup> Kris Hammond, ‘What is Artificial Intelligence?’, Computerworld, 10 April 2015, <https://www.computerworld.com/article/2906336/emergingtech>

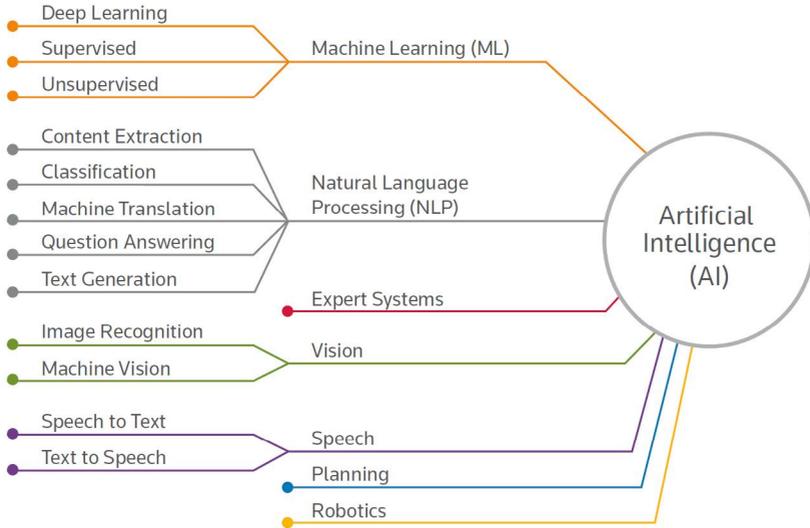
<sup>4</sup> Daniel Ben-Ari et al, ‘Artificial Intelligence in the Practice of Law: An Analysis and Proof of Concept Experiment’ (2017) 23(2) *Richmond Journal of Law & Technology* 2, 6; Ben Allgrove, *Legal Personality for Artificial Intellecs: Pragmatic Solution or Science Fiction?*, Dissertation Submitted for Master of Philosophy (Oxford University, 2015) 3.

<sup>5</sup> Hammond, above n 3; Sean Semmler and Zeeve Rose, ‘Artificial Intelligence: Application Today and Implications Tomorrow’ (2017-18) 16 *Duke Law & Technology Review* 85, 86.

<sup>6</sup> Ben-Ari et al, above n 4, 8.

<sup>7</sup> James A Sherer and Ed Walters, ‘Transitioning from Consumer Tech to Legal Intelligence Engineering’ (2018) *Law Practice* 33.

<sup>8</sup> Sam Adams et al, ‘Mapping the Landscape of Human-Level Artificial General Intelligence’ (2012) 33(1) *AI Magazine* 25, 26.



*Figure 1: Branches of AI<sup>9</sup>*

The goal of Artificial General Intelligence (AGI) is ‘the development and demonstration of systems that exhibit the broad range of general intelligence found in humans.’<sup>10</sup> An example of a test proposed for AGI is Steve Wozniak’s “coffee test”. The test requires a robot to walk into an unfamiliar house and make a cup of coffee.<sup>11</sup> Due to the variability of each house, completing the test would require a number of capabilities that are beyond current state-of-the-art robotics. For example, for the robot to pour coffee from an unfamiliar pot to cup, it will need to use vision to navigate, identify objects, and coordinate fine motor skills. To obtain instructions on making the coffee, it would require speech recognition and natural language processing and generation. As distinct from a narrow AI approach, it would not be feasible to pre-program all of the potential sequences required and the robot will need to find problems and solve them as they arise. The robot would therefore have to demonstrate adaptivity and common sense, and learn by example through machine learning methods.<sup>12</sup> The “coffee test” could easily be completed by most 10-year old humans with little experience<sup>13</sup> and demonstrates the paradox that computers are much better at high level reasoning than low-level sensorimotor tasks.<sup>14</sup>

<sup>9</sup> Michael Mills, ‘Artificial Intelligence in Law: The State of Play’ (Thomson Reuters, 2016) 3.

<sup>10</sup> Adams et al. above n 8.

<sup>11</sup> Ibid 36; see also <https://www.fastcompany.com/1568187/wozniak-could-computer-make-cup-coffee>

<sup>12</sup> Ibid 36.

<sup>13</sup> Ibid 37.

<sup>14</sup> Mark McKamey, ‘Legal Technology: Artificial Intelligence and the Future of Law Practice’ (2017) 22 *Appeal: Review of Current Law and Law Reform* 45.

In terms of legal applications, AI software clearly does not legally reason the way that people do, and arguably it does not need to. Just because computers cannot yet ‘think’, that does not mean they cannot perform some tasks better than humans can.<sup>15</sup> Here, we focus on those AI developments which are increasingly discussed in relation to legal services: expert systems, machine learning and natural language processing.

## HOW DOES AI WORK?

### Expert Systems

Early forays into the field of AI and the law dealt with a branch of AI called ‘expert systems’ (also referred to as knowledge systems<sup>16</sup> or symbolic reasoning systems).<sup>17</sup> Their aim was to capture human expertise and represent it in symbolic form, using logic to program the relationships between different elements.<sup>18</sup> The system must be ‘manually’ programmed, and a human possessing expertise in the relevant domain must be involved in crafting the questions, creating a network of questions so that answers can be generated.<sup>19</sup>

By the 1980s, some legal expert systems were able to provide a certain degree of “legal” reasoning around simple problems.<sup>20</sup> Conditional statement rules (if this, then that) are used to move through a legal problem. In many ways the functioning of a legal expert system, at a base level, is not dissimilar to the way that a lawyer might use an instructions sheet. For relatively simple legal problems, the program guides users by asking a set of questions, in a ‘structured dialog’.<sup>21</sup>

Expert systems were, however, limited in their usefulness due to number of key issues: the variation in legal rules across jurisdictions; imprecise legal concepts such as ‘reasonable foreseeability’; problems when dealing with incomplete information; and the ill-fated proposition that legal knowledge could be reduced to a simple set of rules. This meant that expert systems of the 1980s were only useful for providing a basic level of advice which was often no more useful than a list of boxes to tick.<sup>22</sup>

---

<sup>15</sup> Richard Susskind and Daniel Susskind, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Oxford University Press, 2015).

<sup>16</sup> Richard Gruner, ‘Thinking Like a Lawyer: Expert Systems for Legal Analysis’ (1986) 1(2) *Berkeley Technology Law Journal* 259, 261-62 (describing knowledge systems as simpler, less sophisticated expert systems).

<sup>17</sup> See Kaplan, above n 2, 23; Gruner, above n 16.

<sup>18</sup> Kaplan, above n 2, 23.

<sup>19</sup> Curtis E A Karnow, ‘The Opinion of Machines’ (2017) 19 *Columbia Science and Technology Law Review* 136, 142-43.

<sup>20</sup> Kevin D Ashley, *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age* (Cambridge University Press, 2017) 8-9; Kevin Ashley, ‘Case-Based Reasoning and its Implications for Legal Expert Systems’ (1992) 1(2) *Artificial Intelligence and Law* 113.

<sup>21</sup> Dana Remus and Frank Levy, ‘Can Robots Be Lawyers: Computers, Lawyers, and the Practice of Law’ (2017) 30(3) *Georgetown Journal of Legal Ethics* 501.

<sup>22</sup> Philip Leith, ‘The Rise and Fall of the Legal Expert System’ (2010) 1(1) *European Journal of Law and Technology* 1.

Accordingly, these systems ‘can only be constructed for repetitive and fairly narrow tasks under specific bodies of law’.<sup>23</sup> While expert systems still have application, it has been argued that they will not fundamentally revolutionise the practice of law.<sup>24</sup> Nevertheless these types of system are used today for particular tasks, as they are well-suited to assist a high volume of users with simple legal problems.<sup>25</sup> For this reason, they are becoming popular for online dispute resolution and with legal aid entities.<sup>26</sup>

## Technology and Data

AI appears to be on a rising trajectory,<sup>27</sup> largely due to advancements in technology and the availability of data. First, massive increases in processing power mean that computers can now deal with huge amounts of data. Today, the amount of computer storage available is considerably more than in the early days of AI. Second, there is much more data readily available in electronic form, which means that there is less reliance on a time-consuming process of humans encoding knowledge. The volume of electronic data has been doubling every two years for the last fifty years and is growing exponentially.<sup>28</sup> As a result, there is vastly more data for analysis and learning.

It is important to note, however, that data quantity and quality are critical, as they impact on the functioning of machine learning systems. In fact, poor quality and ‘messy’ data is a substantial impediment to obtaining useful results from a machine learning system, hence the phrase ‘garbage in, garbage out’.<sup>29</sup> Organisations may have volumes of data but not in a useable form, or it may be corrupt, badly labelled, or contain many irrelevancies.<sup>30</sup> Having said this, there are also ways to combat this problem and ‘clean up’ the data.

---

<sup>23</sup> Remus and Levy, above n 21, 525.

<sup>24</sup> Ashley, *Artificial Intelligence and Legal Analytics*, above n 20, 11.

<sup>25</sup> See, eg. <https://www.neotalogic.com/><

<sup>26</sup> See, for example, Joe Tomlinson, *A Primer on the Digitisation of Administrative Tribunals* (University of Sheffield, 12 September 2017) [28]-[33] (describing systems such as SmartSettle, Square Trade and Modria); Daniel W Linna Jnr, ‘Leveraging Technology to Improve Legal Services’ (2017) 96(6) *Michigan Bar Journal* 20 (describing Michigan Help Online and Illinois Legal Aid Online); David Luban, ‘Optimism, Skepticism and Access to Justice’ (2016) *Texas A and M Law Review* 499, 502; Legal Services Corporation, *Report of the Summit on the Use of Technology to Expand Access to Justice* (2013) 1.

<sup>27</sup> ‘Why Artificial Intelligence is Enjoying a Renaissance’, *The Economist*, 15 July 2016, <https://www.economist.com/the-economist-explains/2016/07/15/why-artificial-intelligence-is-enjoying-a-renaissance><

<sup>28</sup> Benjamin Alarie, Anthony Niblett and Albert H Yoon, ‘How Artificial Intelligence Will Affect the Practice of Law’ (2018) 68 Suppl 1 *University of Toronto Law Journal* 106, 124.

<sup>29</sup> Thomas Redman, ‘Data’s Credibility Problem’ (2013) 91(12) *Harvard Business Review* 84; Thomas C Redman, ‘If Your Data is Bad, Your Machine Learning Tools Are Useless’, *Harvard Business Review Digital Articles*, 2 April 2018.

<sup>30</sup> Daniel Shapiro, ‘Artificial Intelligence and Bad Data’, *Towards Data Science*, 6 November 2017, <https://towards-datascience.com/artificial-intelligence-and-bad-data-fbf2564c541a><

## Machine Learning

In recent times, when people talk about artificial intelligence, they are usually referring to machine learning.<sup>31</sup> For lawyers interested in the future of AI and the law, machine learning is key. What takes machine learning beyond expert systems is that the software can adapt itself as it encounters new data and continually enhance its own performance – in other words, “learn” independently.<sup>32</sup>

This means that instead of a person manually inputting a large number of rules, a machine learning system learns through analysing examples.<sup>33</sup> In this way, machine learning software can produce its own “model” and apply it to new and not previously seen data.

Typically, the more data that a machine learning system has to analyse, the greater the accuracy of the model developed.<sup>34</sup> Machine learning systems excel at finding patterns in data:<sup>35</sup>

*Machine learning techniques are concerned with using features or attributes from each example to arrive at the correct label or classification – for example, which key features can be used to distinguish a picture of a cat from a dog? As more examples of cats and dogs are provided, machine learning algorithms can attempt to build models of what underlying distinguishing elements – features – are reliable predictors of whether something is a cat or a dog.<sup>36</sup>*

The model developed is a pattern of statistical relationships that exist between different features of the data in the dataset.<sup>37</sup>

Machine learning software continually refines itself as it encounters new data. Unsurprisingly, it is therefore much better at managing dynamic and complex situations where it would not be possible for a human programmer to predict the best rule to apply.<sup>38</sup> One common example is spam filtering. As a person marks emails as spam (or “unmarks” them by returning them to the inbox) the email program acquires more data to help it predict whether a new incoming email is spam or not.

---

<sup>31</sup> Amir Husain, *The Sentient Machine: The Coming Age of Artificial Intelligence* (Scribner, 2017) 20 (noting that the popular press tends to do this).

<sup>32</sup> Harry Surden, ‘Machine Learning and Law’ (2014) 89(1) *Washington Law Review* 87, 89.

<sup>33</sup> *Ibid* 93.

<sup>34</sup> Matt Kiser, ‘Introduction to Natural Language Processing’, Algorithmia, 11 August 2016, <https://blog.algorithmia.com/introduction-natural-language-processing-nlp/>

<sup>35</sup> Kaplan, above n 2, 27; see also Husain, above n 31, 21.

<sup>36</sup> Husain, above n 31, 21.

<sup>37</sup> Solon Barocas and Andrew D Selbst, ‘Big Data’s Disparate Impact’ (2016) 104 *California Law Review* 671, 677.

<sup>38</sup> Surden, above n 32, 93.

While it might not occur to a human programmer setting up the system to specify that emails originating in Belarus are more likely to be spam,<sup>39</sup> a machine learning system continually modifies itself based on the data it has.

As each spam email arrives, the system ‘adds’ to its model, or factors in the new information. This ‘incremental, adaptive, and iterative process’ can enable a complex and continually developing model.<sup>40</sup> Machine learning is therefore more autonomous than an expert system. It can capture features of phenomena that would not be apparent to a human programmer. While this ability to uncover “new” relationships can be a great strength of machine learning in terms of effectiveness, it can also diminish explainability.

Machine learning may be ‘supervised’ or ‘unsupervised’ and the difference largely depends on the training data for the system. In supervised learning, the data is already labelled (for example, a picture is labelled as a dog or a cat), and the program can therefore identify associations between the data and the labelled outcome, or classification.<sup>41</sup> Most people in fact have participated in ‘training’ a supervised learning system – for example, whenever you are asked to demonstrate that you are not a robot when accessing a website, by identifying pictures (“Select all the pictures which have cats in them”). Each time a person does this, it gives the system more labelled data to work with, helping it to learn which pictures have cats and which do not.

In unsupervised learning, the data is not labelled, and the software searches for patterns in the data. Instead of telling the software which are pictures of cats and which are not, until it learns the difference, it is given enough pictures to discern the pattern itself. Google’s software developers did this with an early version of AlphaGo, which learned to identify pictures of cats by studying thousands of pictures with, and without, cats in them.<sup>42</sup> Although it wasn’t told which pictures had cats and which did not, *‘the system eventually detected the common cat features on its own, and reported its discovery of that common entity’*.<sup>43</sup>

Thus, supervised learning can be described as an exercise in classification; while unsupervised learning is about identifying patterns in data. Supervised learning algorithms are presently the most commonly used in legal tasks.<sup>44</sup>

---

<sup>39</sup> Ibid 94.

<sup>40</sup> Ibid.

<sup>41</sup> Husain, above n 31, 21.

<sup>42</sup> D Silver, et al, ‘Mastering the Game of Go with Deep Neural Networks and Tree Search’ (2016) 529 *Nature* 484.

<sup>43</sup> Karnow, above n 19, 146-47; John Markoff, ‘How Many Computers to Identify a Cat? 16,000’, *New York Times*, 25 June 2012.

<sup>44</sup> Baracos and Selbst, above n 37, 673; David Lehr and Paul Ohm, ‘Playing with the Data: What Legal Scholars Should Learn About Machine Learning’ (2017) 51 *University of California Davis Law Review* 653, 676 (supervised learning algorithms are ‘driving many legally consequential decisions’); Nick Ismail, ‘Artificial intelligence in the legal industry: Adoption and strategy - Part 1’, *Information Age*, 6 August 2018, <https://www.information-age.com/artificial-intelligence-in-the-legal-industry-123473948/>

## Neural networks and deep learning

Current advances in AI relate in large part to progress that has been made in ‘deep learning’. Table 1 summarises the differences between expert systems, machine learning and deep learning. In addition to how each form of AI operates, key elements to be aware of are the relative degree of autonomy that the system is capable of – is it entirely dependent on human programming, or can it develop itself – and the corresponding degree of explainability. Explainability refers to the extent to which humans are able to understand how the system generated its outputs or decisions, from the data that was input.<sup>45</sup>

	Expert systems	Machine learning	Deep learning
Description	Manually programmed system which can guide user through a series of steps to reach pre-determined outcomes	Mathematical analysis of training data, which is used to extrapolate to new data	Sophisticated analysis of data, involving complex calculations and ‘weighting’ of many features
Autonomy	Low, as the program is dependent on human programming	Medium, humans usually guide the program, e.g. by cleaning up data	High, the program decides on feature selection and weighting
Explainability	High	Medium	Low

*Table 1: Summary of AI models<sup>46</sup>*

Artificial neural networks are a kind of machine learning that seeks to replicate the architecture of the human brain. Computational neuroscience hypothesises that the human brain functions via ‘electrochemical activity in networks of brain cells called neurons.’<sup>47</sup> Neurons transmit impulses to one another in a vast network. In seeking to replicate this, artificial neural networks consist of layers, with each layer applying a function to the previous layer. The output of one neuron becomes input for another, so they work together – similarly to the way the human brain is thought to process data, by clustering neurons into groups.<sup>48</sup>

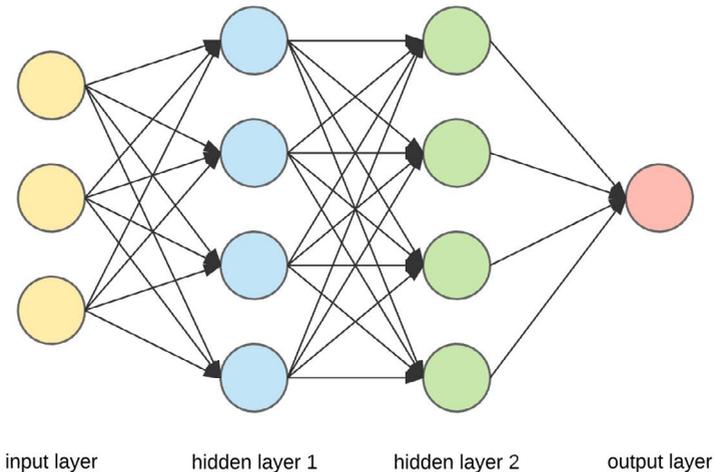
<sup>45</sup> Joshua A Kroll et al, ‘Accountable Algorithms’ (2017) 165 *University of Pennsylvania Law Review* 633.

<sup>46</sup> See also Megan Zweig and Bill Evans, ‘How should the FDA approach the regulation of AI and machine learning in healthcare?’, *Rock Health*, 11 June 2018, <<https://rockhealth.com/how-should-the-fda-approach-the-regulation-of-ai-and-machine-learning-in-healthcare/>>

<sup>47</sup> Russell and Norvig, above n 2, 727.

<sup>48</sup> Daniel Geng and Shannon Shih, ‘Machine Learning Crash Course: Part 3’, *Machine Learning @ Berkeley*, 4 Feb 2017, <<https://ml.berkeley.edu/blog/2017/02/04/tutorial-3/>>

While ‘shallow’ neural networks have been around for a long time,<sup>49</sup> ‘deep’ neural networks are a more recent phenomenon. ‘Deep’ refers to having more than one ‘hidden’ layer (in between input and output layers),<sup>50</sup> as illustrated by Figure 2.



*Figure 2: A deep neural network<sup>51</sup>*

The hidden layers transform the inputs into something that the output layer can use. The output layer transforms the hidden layer activations into the scale specified for the output. A group of data scientists based at UC Berkeley illustrate in greater detail the operation of a neural network using the example of how a system might identify a picture of a dog (see Figure 3). They explain that this is, conceptually, how a neural network functions, once it has been trained with a lot of data.<sup>52</sup>

49 Jürgen Schmidhuber, ‘Deep Learning in Neural Networks: An Overview’ (2015) 61 *Neural Networks* 85, 88.

50 Kaplan, above n 2, 34.

51 Arden Dertat, ‘Applied Deep Learning Part 1: Artificial Neural Networks’, *Towards Data Science*, 8 August 2017, <https://towardsdatascience.com/applied-deep-learning-part-1-artificial-neural-networks-d7834f67a4f6>

52 Geng and Shih, above n 48.

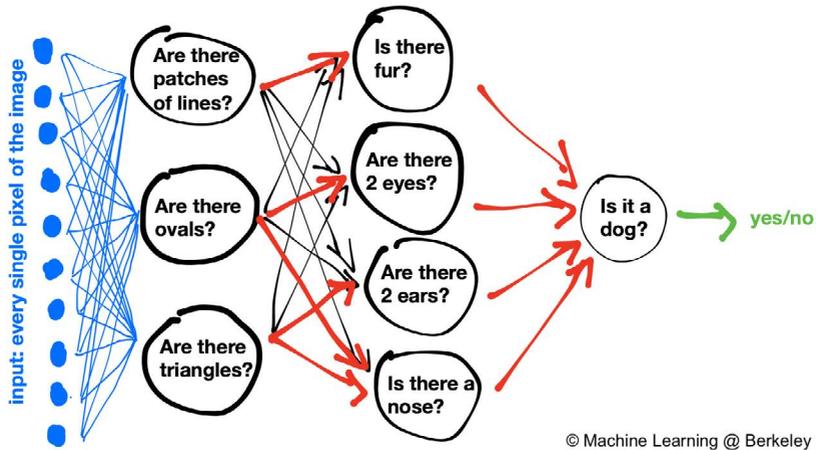


Figure 3: Example of how a neural network works<sup>53</sup>

The hidden layers of neurons and the numerous interactions between neurons reduces explainability as which inputs are relied on, and to what extent, before making a classification is opaque.

Deep learning can be used for supervised or unsupervised learning. So far, most of the benefits of deep learning have involved supervised learning.<sup>54</sup> However, it has enormous potential for unsupervised learning too.

### Natural language processing

Another important development in AI is in the area of text analytics and is referred to as ‘natural language processing’ (NLP). NLP relies on machine learning to analyse patterns of human language as it is used – in other words, for humans and computers to communicate using ‘natural’ language. As language is contextual, NLP uses statistics to work out the probability of words appearing next to one another.<sup>55</sup> The idea that context can be used to establish the meaning of a word is also referred to as ‘semantic analysis’ or ‘latent semantic analysis’.<sup>56</sup>

<sup>53</sup> Ibid.

<sup>54</sup> Andrew Ng, Presentation at Stanford University, 16 July 2014. <https://www.youtube.com/watch?v=W15K9PegQt0>

<sup>55</sup> Benjamin Liu, ‘Can Artificial Intelligence Ever Give Legal Advice?’ (July 2016) *Brief8*, 8.

<sup>56</sup> Remus and Levy, above n 21, 510.

NLP is able to identify the relationship which words in a sentence have to one another, thereby interpreting intent.<sup>57</sup> It is important for law and legal applications because so much of what lawyers do is text-based.

NLP can be used for many language-related tasks, such as automated summarization of text, and can even generate short news articles.<sup>58</sup> Semantic search functions mean that searches, including of legal databases, can be performed using natural language queries instead of Boolean or keyword searches.<sup>59</sup> This has implications for non-lawyer accessibility as well as to potentially improve lawyers' search capacities. Much has also been written about IBM's Watson application. Watson uses both natural language processing and machine learning to answer questions, famously beating champion human competitors at the game show *Jeopardy!* in 2011. Since then, a legal version of Watson called ROSS Intelligence has been developed.<sup>60</sup> ROSS uses the same technology to deliver a legal search engine, where the user can enter his or her question in natural language and does not have to engage with Boolean searching. It is claimed that in this way, ROSS provides a superior and more intuitive means of searching.<sup>61</sup> In addition, rather than retrieve an entire case, ROSS retrieves the relevant sections of the case(s).<sup>62</sup>

---

57 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, 447.

58 Tim Adams, 'And the Pulitzer goes to... a computer', *The Guardian*, 28 June 2015, <https://www.theguardian.com/technology/2015/jun/28/computer-writing-journalism-artificial-intelligence> (Kiser, above n 34).

59 John O' McGinnis and Russell G Pearce, 'The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers in the Delivery of Legal Services' (2014) 82(6) *Fordham Law Review* 3041.

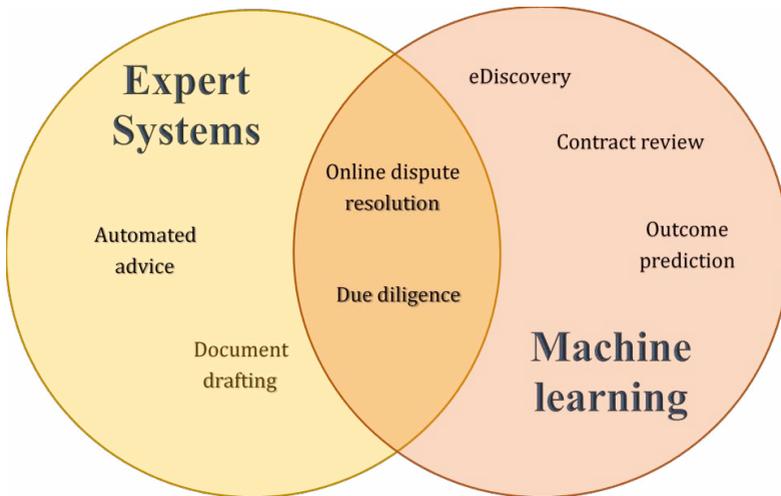
60 <https://rossintelligence.com/>

61 See [Blue Hill Research, ROSS Intelligence and Artificial Intelligence in Legal Research \(Feb 2017\)](#) 5.

62 Remus and Levy, above n 21, 521; Andrew Arruda quoted in Kim-Mai Cutler, 'YC's ROSS Intelligence Leverages IBM's Watson to Make Sense of Legal Knowledge', 28 July 2015, <https://techcrunch.com/2015/07/27/ross-intelligence/>. See also Kevin Van Paassen, 'University of Toronto's next lawyer: A computer program named Ross', *The Globe and Mail*, 11 December 2014, updated 25 March 2017.

## HOW IS AI USED IN LEGAL SERVICES?

One suggested definition of the use of AI in the practice of law is ‘the theory and development of processes performed by software instead of a legal practitioner, whose outcome is the same as if a legal practitioner had done the work’.<sup>63</sup> AI is currently being used in legal services in different ways, including through the mechanisation of legal services (including technology assisted review of documents, and legal searches), legal document creation, and – perhaps most controversially – prediction of legal outcomes. Two US commentators have argued that AI is best positioned to do two things – take over ‘repetitive’ tasks like discovery; and produce ‘commodity legal documents’.<sup>64</sup>



*Figure 4: Legal applications of AI*

<sup>63</sup> Sergio David Becerra, 'The Rise of Artificial Intelligence in the Legal Field: Where We Are and Where We Are Going' (2018) 11 *Journal of Business, Entrepreneurship & Law* 27, 38.

<sup>64</sup> Dennis M Horn and Ira Meislik, 'How to Ride the Coming Tidal Wave of Technology and Competition' (2018) 32(6) *Probate and Property* 9.

## eDiscovery

Perhaps the best known application of machine learning to law is Technology Assisted Review (TAR), also referred to as ‘computer-assisted review’,<sup>65</sup> which can be used in discovery<sup>66</sup> where the material to be processed is voluminous. Given the volume of electronically stored information (ESI) in present day large-scale litigation, it is a means of automating the review of electronic documents. TAR has been used in US courts since 2012,<sup>67</sup> in English and Irish court proceedings<sup>68</sup> and was approved by the Supreme Court of Victoria in 2016, in *McConnell Dowell Constructors v Santam*.<sup>69</sup> Following the decision in *McConnell Dowell*, the Supreme Court of Victoria introduced a new Practice Note dealing with technology in civil litigation that endorsed TAR in larger cases as ordinarily being an accepted method of conducting a reasonable search. Other Australian jurisdictions have employed TAR, although there has not been judicial discussion of its application.<sup>70</sup>

It is now generally accepted that when a high number of documents are involved, TAR is likely to be considerably faster (and therefore more cost-effective) and more accurate, than a traditional ‘manual’ or ‘linear’ review of documents for discovery,<sup>71</sup> even where keyword searches are used.<sup>72</sup> TAR works by using machine learning’s capacity to identify patterns in data, including in textual data. It is an example of supervised machine learning: the program is provided with a set of documents referred to as a ‘seed set’ (also referred to as a starter set or training set). The seed set may be randomly compiled from all the potentially discoverable documents, or documents may be selected for particular characteristics.

---

<sup>65</sup> Further terms include ‘computer-aided review’, ‘predictive coding’ and ‘content-based advanced analytics’: Maura R Grossman and Gordon V Cormack, ‘The Grossman-Cormack Glossary of Technology-Assisted Review’ (2014) 7(1) *Federal Courts Law Review* 85.

<sup>66</sup> Discovery is the production of documents to the opposing side as part of the trial process.

<sup>67</sup> 287 ER D 182 (SDNY 2012). The decision was upheld on appeal. For a summary of US cases see Julia L Brickell and Peter J Pizzi, ‘Towards a Synthesis of Judicial Perspectives on Technology-Assisted Review’ (2015) 82 *Defence Counsel Journal* 309; and Ernst and Young, *Insiders’ Guide to Technology-Assisted Review* (Wiley, 2015) ‘Ch 2: TAR and the Case Law’.

<sup>68</sup> *Irish Bank Resolution Corporation Limited v Quinn* [2015] IEHC 175 and *Pyrrho Investments Ltd v MWB Property Ltd* [2016] EWHC 256 (Ch).

<sup>69</sup> *McConnell Dowell Constructors (Aust) Pty Ltd v Santam Ltd and Others (No 1)* [2016] VSC 734 (Vickery J) (“*McConnell Dowell*”).

<sup>70</sup> See eg *Cantor v Audi Australia Pty Limited (No 3)* [2017] FCA 1079 where Foster J in the Federal Court put off an application by the ACCC to appoint a referee to consider the question of whether TAR should be ordered; Practice Note SC Gen 7 of the NSW Supreme Court which is sufficiently broad to enable TAR: Michael Legg and Thomas Davey, ‘Predictive Coding: Machine Learning Disrupts Discovery’ (2017) 32 *Law Society Journal* 82.

<sup>71</sup> See, eg, Karnow, above n 19, 141.

<sup>72</sup> Maura R Grossman and Gordon V Cormack, ‘Technology-Assisted Review in E-Discovery Can be More Effective and More Efficient than Exhaustive Manual Review’ (2011) 17 *Richmond Journal of Law & Technology* 1; Bennett B Borden and Jason R Baron, ‘Finding the Signal in the Noise: Information Governance, Analytics, and the Future of Legal Practice’ (2014) 20 *Richmond Journal of Law & Technology* 1, 7, 16; Ernst and Young, above n 67, ‘Ch 3: The Economics of TAR’.

A human (lawyer) reviewer then codes the documents in the seed set, labelling them (for example) as ‘relevant, not relevant, privileged, or not privileged’.<sup>73</sup> Using this information, the program applies this to other documents.<sup>74</sup> Just as a program can eventually successfully attach a label to a not-previously seen picture of a cat, once trained it can also successfully identify which documents in the discovery are relevant, and which are not, ‘with a high degree of accuracy’.<sup>75</sup> Thus, from the seed set the software creates ‘a predictive model, a kind of profile’<sup>76</sup> of the different types of documents, and this ‘mathematical model... can then predict the classifications of other documents in that dataset’.<sup>77</sup> Ultimately, the program generates a probability that a particular item is relevant/not relevant.

### Document review and due diligence

Machine learning can be used also for review of specific types of document, such as contracts.<sup>78</sup> After being trained through being provided with many examples, the software can identify (for instance) different types of contracts by using pattern recognition. This has two main benefits: firstly, it can be used as an organisational tool by companies, as it enables sophisticated contract management. That is, the program can automatically extract information such as particular clauses, or parties’ names, and so on.<sup>79</sup> It can, for example, flag all the contracts which are to expire within a certain time period or identify all of those which contain a certain clause.

Secondly, programs can review contracts of the particular type they are trained in, and identify the wording of clauses.<sup>80</sup> Contracts which are largely standardised are susceptible to the automation of review – such as commercial leases or non-disclosure agreements.<sup>81</sup> One company analysed 250,000 employment agreements to build a program that can, when presented with a new contract, identify which clauses are non-standard.<sup>82</sup> The program does not, of course, explain what should be done about the clause, or its meaning. The founder of one start-up has explained: ‘A person using our system misses less than they would otherwise. It’s an enhancement tool’.<sup>83</sup>

73 Matthew Paulbeck, ‘The Ethics of Predictive Coding: Transparency and Judgment-Formed Seed Sets’ (2017) 30(4) *Georgetown Journal of Legal Ethics* 971.

74 For a comprehensive description, see Bolch Judicial Institute, ‘Technology Assisted Review (TAR) Guidelines’, January 2019 (Duke Law School) <https://www.edrm.net/wp-content/uploads/2019/02/TAR-Guidelines-Final.pdf>

75 *Ibid.*

76 Ashley, above n 20, 241.

77 Shannon Brown, ‘Peeking Inside the Black Box: A Preliminary Survey of Technology Assisted Review (TAR) and Predictive Coding Algorithms for Ediscovery’ (2016) 21 *Suffolk Journal of Trial & Appellate Advocacy* 221, [2.1] (see generally for a comprehensive technical overview of the TAR process).

78 Eg. Kira Systems, COIN; Otto Hanson, ‘Product Review: Kira – Contract Extraction Software for M&A Due Diligence’ (2018) 47 *Colorado Lawyer* 13.

79 Beverly Rich, ‘How AI Is Changing Contracts’, *Harvard Business Review*, 12 February 2018, 3.

80 Sherer and Walters, above n 7, 36; Semmler and Rose, above n 5, 89.

81 See, eg. LawGeex ([www.lawgeex.com](http://www.lawgeex.com)).

82 K M Standards (<http://kmstandards.com/>).

83 Victor Li, ‘Technology rewires the drafting and reviewing of contracts’, *ABA Journal*, November 2014 (quoting Noah Waisberg).

Prior to entering into a transaction with or for a company, the risks associated with the entity must be assessed, which generally means reviewing the company's contracts, among other things. Traditionally, this work has been done by teams of lawyers reading documents.<sup>84</sup> However, there are now systems which, using AI, can also perform some of this work effectively. Due diligence differs from discovery because it has an exploratory element as well as a 'structured component'.<sup>85</sup> Various software companies have developed products which can work on the 'structured component'.<sup>86</sup> This includes programs which can identify different types of documents and classify them. A review of one such program explains: '[It] has more than 430 provisions that it is already trained to look for, but it also employs artificial intelligence to let you train [it] to look for more unique provisions specific to your deal type'.<sup>87</sup> Like other applications of machine learning, the software will continue to learn as it is exposed to more documents. But it has been argued that this type of software cannot yet do a good job of the 'unstructured' component of due diligence – recognizing that something does not "fit" or is not quite right.<sup>88</sup>

Nevertheless, there are a number of companies now offering legal document review services.<sup>89</sup> The biggest law firms have been entering into agreements with these software providers in order to streamline the firms' services. It's claimed that these programs remove the need for lawyers to undertake the initial, time-consuming review of documents and therefore focus only on problematic clauses or elements flagged by the program. These types of program require scale – the more similar documents they are trained with, the greater ability there will be to 'recognise' elements in new documents. They can also be customised – for example, Kira Systems offers software which is 'trained' on various types of agreement, but which can be further trained using a firm's own documents.<sup>90</sup>

---

<sup>84</sup> Hanson, above n 77, 13.

<sup>85</sup> Remus and Levy, above n 21, 21.

<sup>86</sup> *Ibid.*

<sup>87</sup> Hanson, above n 77, 13.

<sup>88</sup> Remus and Levy, above n 21, 21-22.

<sup>89</sup> For example, Luminance (<https://www.luminance.com/>), Kira Systems (<https://kirasystems.com/>) and iManage (<https://immanage.com/product/artificial-intelligence/>).

<sup>90</sup> 'Kira is ready to use out-of-the-box and can be highly customized. The software comes pre-trained with over 50 real estate lease provision models... Customers have also used Kira to successfully train their own lease provision models': <https://kirasystems.com/solutions/lease-abstraction/>

## Document drafting and automated advice

The use of software to automate the drafting of legal forms and documents has been around for some time.<sup>91</sup> Companies such as Desktop Lawyer,<sup>92</sup> LegalZoom<sup>93</sup> and Rocket Lawyer<sup>94</sup> use expert systems to enable lay users to carry out straightforward legal tasks like incorporating a company or drawing up a will, by answering a series of questions.<sup>95</sup> Forms, letters or agreements are populated with the answers provided by the consumer or client. Most services offered direct to consumers provide non-contentious, simple documents which tend to be standardised. There are now hundreds of websites offering everything from advance care directives, to patent applications, to challenging parking fines. In Australia, LegalVision, LawPath and ClickLaw offer similar services.

There has also been huge growth in legal apps and chatbots. The parking fines challenges chatbot, DoNotPay,<sup>96</sup> credited with being the first ‘legal chatbot’, asks the user various questions using NLP and generates a simple letter challenging the fine using the responses.<sup>97</sup> Norton Rose Fulbright recently announced the creation of a similar ‘chatbot’ to assist its clients in managing their obligations under new Australian privacy laws.<sup>98</sup> Legal practices can also build their own chatbots using off the shelf products.<sup>99</sup>

For lawyers, using document automation software is a step up from using templates or creating documents based on previous iterations. Most document automation software is centrally hosted and available to lawyers or legal practices on a subscription basis. Lawyers can choose whether to automate their own existing documents (the firm’s existing precedents); or subscribe to a system which is prepopulated with documents designed by the service.

---

91 Kathryn D Betts and Kyle R Jaep, ‘The Dawn of Fully Automated Contract Drafting: Machine Learning Breathes New Life into a Decades-Old Promise’ (2017) 15 *Duke Law & Technology Review* 216, 218; citing Kenneth I Guthrie, ‘Document Assembly Software Systems’ (1995) 9 *Probate and Property* 26, 27. See also Catherine J Lanctot, ‘Regulating Legal Advice in Cyberspace’ (2002) 16 *St. John’s Journal of Legal Commentary* 569, 579.

92 <https://desktoplawyer.secureclient.co.uk/dtl/>

93 [www.legalzoom.com](http://www.legalzoom.com)

94 [www.rocketlawyer.com](http://www.rocketlawyer.com)

95 See Gerard J Clark, ‘Internet Wars: The Bar against the Websites’ (2013) 13 *Journal of High Technology Law* 247.

96 Elena Cresci, ‘Chatbot that overturned 160,000 parking fines now helping refugees claim asylum’, *The Guardian online*, 6 March 2017, <https://www.theguardian.com/technology/2017/mar/06/chatbot-donotpay-refugees-claim-asylum-legal-aid>

97 Shannon Liao, ‘Chatbot lets you sue Equifax for up to \$25,000 without a lawyer’, *The Verge*, 11 September 2017.

98 Press Release, ‘Norton Rose Fulbright launches first Australian law firm chatbot to help manage data breach’, 13 December 2017, <http://www.nortonrosefulbright.com/news/159704/norton-rose-fulbright-launches-first-australian-law-firm-chatbot-to-help-manage-data-breach> See also <http://www.nortonrosefulbright.com/knowledge/publications/147579/australian-privacy-compliance-packages>

99 Eg Josef (<https://joseflegal.com>); Neota

This kind of process could also be used by legal practices to streamline back office tasks: for example, by prompting staff to step through a process; or for client intake, as the potential client can provide key information prior to a first meeting. Outcome prediction

Increasing volumes of data available feed ever more sophisticated analyses of that data. One application of this is to use machine learning to make predictions about people's future behaviour. While this is one of its more controversial applications, it occurs routinely in some areas: AI is used in the financial services industry to assign credit scores and approve loans;<sup>100</sup> and in the legal arena, in policing,<sup>101</sup> to assess the likelihood that a person will skip bail<sup>102</sup> or go on to commit further crimes if paroled.<sup>103</sup> Concerns have been voiced about the use of AI software to make such predictions, including whether this can amount to a denial of due process/procedural fairness.<sup>104</sup>

It is argued that 'big data' can bring a more fundamental change to the nature of legal work, 'by looking to statistical patterns, predictors, and correlations, in addition to the legal rules that purportedly control outcomes'.<sup>105</sup> As with the document review applications discussed above, this works effectively for specific types of application which are narrowly framed and for which there is a high volume of data. AI's ability to detect patterns in large volumes of data, including in unstructured and unlabelled data, also enables (for example) sophisticated fraud detection; and corresponding management of risk and compliance. A system can detect, often in real time, anomalous patterns of behaviour (such as credit card spending) which can then be escalated for human investigation.

---

100 Mikella Hurley and Julius Adebayo, 'Credit Scoring in the Era of Big Data' (2016) 18 *Yale Journal of Law and Technology* 148.

101 Justin Jouvenal, 'Police are using software to predict crime. Is it a 'holy grail' or biased against minorities?', *Washington Post*, 17 November 2016.

102 Jon Kleinberg, et al, 'Human Decisions and Machine Predictions' (2017) Working Paper No. 23180, National Bureau of Economic Research.

103 See the examples given by Kroll et al, above n 45, 636.

104 Julia Angwin et al, 'Machine Bias', *ProPublica*, 23 May 2016, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (arguing that software used for undertaking risk assessments is biased against racial minorities); Lindsey Barrett, 'Reasonably Suspicious Algorithms: Predictive Policing at the United States Border' (2017) 41(3) *New York University Review of Law & Social Change* 327-366; Recent Cases, 'Wisconsin Supreme Court Requires Warning before Use of Algorithmic Risk Assessments in Sentencing' (2017) 130 *Harvard Law Review* 1530. Others argue that AI has potential for the administrative state and is not necessarily negative: Cary Coglianese and David Lehr, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 *Georgetown Law Journal* 1147.

105 Dru Stevenson and Nicholas J Wagoner, 'Bargaining in the Shadow of Big Data' (2015) 67 *Florida Law Review* 1337, 1342.

Another way that AI might be used is to predict the outcome of a stated case by reference to earlier decisions.<sup>106</sup> At the moment, the data used to analyse and predict case outcomes is based on factors external to a case (such as court location, judge, lawyers and so on) rather than text itself. But there have been attempts, within subject-specific domains, to utilise the predictive powers of AI software and there are now a number of providers who are reputedly doing so.<sup>107</sup> For example, London-based CaseCrunch claimed that its software was more accurate than lawyers in predicting the outcome of decisions made by the Financial Ombudsman on payment protection insurance, though one criticism was that the lawyers were not specialists in the area.<sup>108</sup>

In a 2016 article, a group of European computer scientists described their attempts to ‘predict’ the outcome of cases heard by the European Court of Human Rights using machine learning.<sup>109</sup> Other scholars have argued that the type of prediction which is possible from machine learning algorithms can simply never replace legal reasoning.<sup>110</sup> Yet to an extent this is already occurring – at least in so far as it enables lawyers to improve their research, and gives additional knowledge about courts, judges and persuasiveness of different motions. This type of software could be used by lawyers to assist in advising client about the merits of their case; by litigation funders, to determine whether funding of a particular lawsuit is a good investment;<sup>111</sup> or by Legal Aid bodies to determine whether a person should receive legal aid assistance.

In terms of legal practice, it is argued that though big data is garnering a lot of attention, the first step for law firms is actually to analyse their own “small data”.<sup>112</sup> In other words, using existing data resources – about clients, similar case outcomes and settlements – is the first step to increasing practice efficiency.<sup>113</sup>

106 See J Dixon, ‘Review of Legal Analytics Platform,’ *Litigation World*, 23 September 2016.

107 Eg. Ravel Law (including Judge Analytics) (<https://home.ravelaw.com/>), Lex Machina (<https://lexmachina.com/>).  
108 The CaseCrunch website claims that ‘lawyers scored an accuracy of 62.3% while the software scored 86.6%. No detail as to how this was calculated is provided. See also [Rory Cellan-Jones, ‘The robot lawyers are here – and they’re winning,’ BBC News online, 1 November 2017](#); [Jason Tashea, ‘Artificial intelligence software outperforms lawyers \(without subject matter expertise\) in matchup,’ ABA Journal, 3 November 2017](#).

109 Nikolaos Aletras et al, ‘Predicting Judicial Decisions of the European Court of Human Rights: A Natural Language Processing Perspective’ (2016) 2 *PeerJ Computer Science* 92, <https://doi.org/10.7717/peerj-cs.93>; see also <http://www.ucl.ac.uk/news/news-articles/1016/241016-AI-predicts-outcomes-human-rights-trials>. For a critique of the methods used by Aletras et al, see Frank Pasquale and Glyn Cashwell, ‘Prediction, Persuasion, and the Jurisprudence of Behaviourism’ (2018) 68: Suppl 1 *University of Toronto Law Journal* 63.

110 Cass R Sunstein, ‘Of Artificial Intelligence and Legal Reasoning’ (2001) 8 *University of Chicago Law School Roundtable* 29, 32–34 (computer programs do not reason analogically the way humans do).

111 See, eg, Joshua Hunt, ‘What Litigation Finance is Really About’, *The New Yorker*, 1 September 2016, <https://www.newyorker.com/business/currency/what-litigation-finance-is-really-about>.

112 Ed Walters and Morgan Morrisette Wright, ‘The Decision Advantage: Making Small Data Work for Your Firm’ (2018) 89 *Oklahoma Bar Journal* 32.

113 *Ibid.*

## Online dispute resolution

Online dispute resolution (ODR) may refer both to alternative dispute resolution which is conducted online, and to systems of online courts.<sup>114</sup> While ODR may just involve traditional ADR processes which are conducted online or through electronic means, it may be qualitatively different, and ‘replace or significantly reduce the role of humans and instead use advanced artificial intelligence (including algorithms, machine learning and big data) to become the third party that performs the mediation or decision making’.<sup>115</sup> ODR and online courts are both advanced as capable of promoting access to justice.<sup>116</sup> This is for reasons of cost-effectiveness, but also accessibility for users who are located in remote areas, have disability, and so on.

Recently, the former Chief Justice of the Supreme Court of Victoria, Professor Marilyn Warren, suggested that dispute resolution based on algorithmic decision making could be used to deal with small civil claims.<sup>117</sup> By analysing previous decisions, the software can reportedly advise parties of the probability of success and the range of outcomes. The negative side of this is the same criticism made of using past decisions to make ‘predictions’ about any likely outcome. It may prevent someone with an unlikely-to-succeed but valid claim from bringing the claim and cause the law to stagnate.

---

114 Michael Legg, ‘The Future of Dispute Resolution: Online ADR and Online Courts’ (2016) 27 *Australasian Dispute Resolution Journal* 277.

115 *Ibid.*

116 Ayelet Sela, ‘Streamlining Justice: How Online Courts Can Resolve the Challenges of Pro Se Litigation’ (2017) 26 *Cornell Journal of Law & Public Policy* 331.

117 Monash University, ‘AI may decide the outcome in civil disputes’, 14 March 2018, <<https://lens.monash.edu/2018/03/14/1331365/ai-may-decide-civil-disputes>>

## ABOUT THE AUTHORS

*The FLIP Stream is primarily conducted by Professor Michael Legg, Dr Justine Rogers and Dr Felicity Bell.*



### **PROFESSOR MICHAEL LEGG, DIRECTOR**

Professor Michael Legg's research interests are in dispute resolution, access to justice and the legal profession. He has previously written on the use of technology assisted review in litigation and online dispute resolution / courts.

He was the Chair of the UNSW Law School's technology curriculum review which examined the ramifications of the impact of technology on the legal profession for legal education.

In 2017 he was awarded Academic of the Year at the *Lawyers Weekly* Australian Law Awards for his innovative teaching of technology and legal practice, especially in relation to litigation and alternative dispute resolution, and engagement with the legal profession. In 2016 he received the Dean's Award for Impact and Engagement.

Michael is admitted to practice in the Supreme Court of NSW, Federal Court of Australia, High Court of Australia and in the State and Federal courts of New York. He holds law degrees from UNSW and the University of California, Berkeley.

Michael is a member of the Law Society of New South Wales' Future Committee, the Law Council of Australia's Class Actions Committee and a Board Member of the Australian Pro Bono Centre.

---



### **DR JUSTINE ROGERS, DEPUTY DIRECTOR**

Dr Justine Rogers researches and teaches in professions, professional work and professional ethics. Her research examines how the changing nature of professions raises urgent global challenges - but also possibilities - for issues of ethics, identities, expertise, and ultimately the public good. From 2014-2018, Justine was a Chief Investigator of an ARC linkage grant with the Professional Standards Council on the future of the profession.

Justine is also convenor of UNSW Law's core UG and JD applied ethics course, which she was appointed in 2013 to design. Her teaching innovations, centred on group-based deliberative ethics, have been recognised and replicated nationally and internationally. Justine was an Academic of the Year Finalist (2016) in the Annual Australian Law Awards, and Women Legal Academic of the Year Finalist (2016). In addition, her course is used at UNSW as an exemplar of blended learning.

Justine has consulted on ethical culture and infrastructure to the legal profession, law firms and a major bank. She has been invited to write on ethics and its meanings in a range of different legal, financial and medical contexts.

---



### **DR. FELICITY BELL, RESEARCH FELLOW**

Dr Felicity Bell holds a BA/LLB (Hons I) from the University of Melbourne and a PhD from the University of Sydney. Felicity has researched and published in family law, children and the law, and examines issues relating to best practice, identity and conceptions of ethics among lawyers.

She has worked with NSW Legal Aid on facilitating children's participation in family law processes. Felicity has taught family law, legal professional ethics and property law, most recently working as a lecturer at the University of Wollongong where she was also the lawyer member of the Human Research Ethics Committee.



# flip stream

THE LAW SOCIETY OF NSW  
IN COLLABORATION WITH UNSW LAW

THE LAW SOCIETY OF NEW SOUTH WALES

170 Phillip Street, Sydney NSW 2000, DX 362 Sydney

ACN 000 000 699 ABN 98 696 304 966

T +61 2 9926 0333

[lawsociety.com.au](http://lawsociety.com.au)