Submission for Australian Academy of Law Essay Competition

Topic

What effect have the advances in technology (including in artificial intelligence) had upon the discipline of law in academia, the practising profession and the courts, and how may that effect change over the next ten years? What steps should be taken now to harness the benefits and limit the detriments of those advances?

Author

Lyria Bennett Moses

UNSW Australia

Faculty of Law

Building F8

Sydney NSW 2052

lyria@unsw.edu.au

(02) 9385 2254

Abstract

Advances in technology, in particular in artificial intelligence, will continue to have a significant impact on the discipline of law in academia, the practicing profession and the courts. While technological forecasting is a dangerous game, current trends suggest that over the next ten years there will likely be greater reliance on data analytic tools in assessing students, predicting judicial outcomes and making decisions about criminal defendants both pre- and post-conviction. There is also likely to be greater diffusion of expert systems offering standardised legal advice and legal documents, although it is less likely that there will be significant technological innovation in that field.

There are significant differences between an artificial intelligence that mirrors doctrinal logic (expert systems) and an artificial intelligence based on projection from empirical observation (data analytics). In particular, few legal professionals understand the mechanisms through which data analytics produces predictions. The limitations inherent and assumptions embedded in these tools are thus often poorly understood by those using them.

This essay will explore the limitations of artificial intelligence technologies by considering the ways in which what they produce (for clients, law students and society) differs from what they replace. Ultimately, if we, as legal professionals, want to harness the benefits and limit the detriments of new artificial intelligence technologies, we need to understand what their limitations are, what assumptions are embedded within them and how they might undermine appropriate decision-making in legal practice, legal academia and, most crucially, the judiciary.

1. Introduction

Law is often perceived as being non-technological, even old fashioned. For example, an article in the Australian Financial Review about the growth in online legal services is accompanied by a photograph of a barrister's wig sitting adjacent to a shelf of bound copies of the Australian Law Reports.[1] However, technology has always had a significant role on the practice of law, the operation of courts and legal academia. Successive communications technologies (the postal system, telephones, facsimiles, email) have increased the speed and decreased the cost of communicating with clients and other parties. Word processing has changed how law students write essays and contributed to substantial increases in the length of judgments. Case management software for law firms, in house legal teams and courts has increased in sophistication and automation.[2] Firms are providing more sophisticated online platforms through which to recruit new clients.[3] There are digital marketplaces for legal services that can connect those seeking legal advice, or law firms managing excess capacity, to lawyers able and willing to work.[4] Legal information is made more easily available through online platforms including databases such as Austlii, commercial providers such as LexisNexis, and information sheets created by law firms, government departments and the not for profit sector. Automated online citators, such as LawCite, are a significant improvement on stamps and stickers once painstakingly affixed to law report volumes.

---

[1] Marianna Papadakis, 'LawPath pursues growth in online legal services', *Australian Financial Review* (9 May 2014).

[2] See, eg, Riverview Law <http://www.riverviewlaw.com>; Sylvia Kriven, 'New IT System to transform justice' (2014) 36 *Bulletin (Law Society of SA)* 9.

[3] For example, Legal Vision which targets small and medium businesses <https://legalvision.com.au> and Axiom <http://axiomlaw.com> which assists in house teams to 'deliver more efficient and effective legal support and improve legal processes'.

[4] For example, see LawPath <https://lawpath.com.au/quick-quotes> and Crowd & Co <https://crowdandco.com.au>.

People can post legal questions on online fora, hoping that an answer will emerge from legal professionals using the site.[5] Information on specific legal issues can also be accessed through legal expert systems, prepared in advance by someone with the relevant expertise to give answers to a range of pre-conceived situations based on responses to pre-written questions.[6] A similar technology also enables automated document assembly, where tailored legal documents have terms set through online interrogation.[7] Litigation can be minimised through the use of online dispute resolution platforms.[8] Technology can even replace law in some circumstances by directly constraining or encouraging action, as in the case of speed humps, computer code and 'privacy by design'.[9]

In legal academia, the use of technologies both within and outside classrooms is pervasive, with most courses having an on-line presence through platforms such as Moodle and Blackboard. Fully online courses are also available, either as a course open to enrolled students, or through a MOOC (massive open online course) platform. While the technologies driving online learning are improving over time, the concept is not new. Already in 1998, Graham Greenleaf taught 'Information Technology Law' at the University of New South Wales remotely, using a basic web platform.

---

[5] See, for example, LawAnswers <http://www.lawanswers.com.au>.

[6] See, for example, using the Neota Logic platform <http://www.neotalogic.com.au>. For example, Justice Connect use this platform to provide advice on establishing a not for profit, see Not-for-profit Law Information Hub <http://www.nfplaw.org.au/gettingstarted>.

[7] Thomson Reuters Cleardocs <www.cleardocs.com>; Law Central Online <http://www.lawcentral.com.au>; legalzoom <http://www.legalzoom.com>; LawPath <http://lawpath.com.au>.

[8] For consumer disputes, see Modria <http://modria.com>. For a broader dispute base, see Rechtwijzer 2.0 <http://www.hiil.org/project/rechtwijzer>. For the use of technology in larger scale innovations in access to justice, see generally Innovating Justice <www.hiil.org>.

[9] See generally Lawrence Lessig, *Code and Other Laws of Cyberspace* (Basic Books 1999).

Examining the effect of technology as a *general* category on lawyers working in a range of roles would be an enormous task. Law is embedded in technologies, whether the technologies of writing and printing, technologies of accessing legal information, or directly in computer code.[10] Many diverse software applications will have a significant impact on professional work, professional incomes and access to legal services. Rather than accounting for broad-ranging impacts, this essay will focus on a category of technology that is particularly problematic, primarily because it is rarely understood by lawyers – artificial intelligence.

Thus this essay will explore the particular questions raised by the use of artificial intelligence in the delivery of legal services and legal education, as well as in courts and judicial decision-making. After explaining what is meant by artificial intelligence and distinguishing different types of artificial intelligence (Part 2), the essay will describe how these have been applied in legal practice, legal academia and courts (Part 3). Part 4 of the essay explores the limitations of data driven artificial intelligence from the perspective of clients, students and society, making some predictions and leading to a conclusion (Part 5) on the importance of understanding for harnessing the benefits and limiting the detriments of such applications.

2. Artificial intelligences

In order to evaluate the effect and impact of a technology such as artificial intelligence, it is necessary to understand what it *is*. There are different tests for artificial intelligence, including the famous Turing test, according to which a computer demonstrates intelligence if

[10] Mireille Hildebrandt, 'A Vision of Ambient Law' in Roger Brownsword and Karen Yeung (eds), *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes* (Hart, 2008); Lessig, above n 9.

a human who interrogates both that computer and another human is unable to determine which is the machine. [11] However, the term ‘artificial intelligence’ is commonly used more broadly to describe a machine performing tasks that would ordinarily require human intelligence. The latter definition is adopted here.

One kind of artificial intelligence technology used in law, the expert system, is designed to mimic the thought process of a human expert.[12] Expert systems rely on expertise drawn from the minds of human legal experts on a particular topic, either working with a ‘knowledge engineer’, relying on their own system design skills or deploying a software tool such as Neota Logic. They can be used in legal practice and legal academia. An expert system can ask the user questions to draw out a factual matrix or assess a student’s knowledge. The expert system relies on the expertise with which it is programmed to provide advice or feedback. The conclusions of the system can be accompanied by reasons in the form of a series of logical statements and citations to relevant sources of authority.

These systems can be useful, and have a range of current applications, discussed in Part 3 below. However, they also have limits. A legal expert system must be constantly updated – new legislation or cases that impact on legal advice or course content will need to be programmed into the system, requiring a similar level of expertise to that deployed in the original design. A legal expert system is also limited by foresight – it is unlikely that every situation and variable within the system’s scope can be foreseen. Further, when giving legal advice, there is often a requirement for appreciation of subtle circumstantial differences to assess whether conduct meets a generally crafted standard such as ‘reasonableness’. A person designing a legal expert system is left with a choice – either ask the client directly whether

[11] Alan Turing, ‘Computing Machinery and Intelligence’ (1950) 59(236) *Mind* 433.

[12] Alan Tyree, *Expert Systems in Law* (Prentice Hall, 1989) 1.

particular conduct was ‘reasonable’ (providing them with a description of the legal standard) or attempt to craft questions that will elicit all of the information that might impact on an assessment of ‘reasonableness’. The first task relies on the client for something better assessed by a lawyer, the second is practically impossible. Thus while expert systems can deal with simple scenarios (such as determining whether a person is eligible for a government program based on fixed statutory criteria), it will perform poorly at giving more complex legal advice.

Expert systems are intelligent to the extent that they can mirror the human expertise of their creator. They assume that legal questions are answered by ‘thinking like a lawyer’; doctrinal reasoning remains at the core.

A very different way to simulate intelligence is through the deployment of data analytic techniques such as machine learning. With machine learning, one does not need expertise or knowledge *ex ante* provided one has a sufficient quantity of sample data. The computer ‘learns’ from correlations and patterns in historic data. This can be done through supervised learning, where a human provides the machine with outputs for a set of inputs in a ‘training set’, or through unsupervised learning, where the computer alone identifies clusters and regularities within data. Once correlations, patterns or clusters are identified, the computer is able to classify new inputs (for example, into high and low risk categories), provide outputs for new inputs (for example, risk scores or estimated damages) or place new data into established clusters (for example, different types of offenders or documents). Machine learning acts intelligently in the sense that it learns over time, and can thus be responsive to feedback, and in the sense that the patterns learnt yield useful predictions or insights. However, machine learning does not follow the same logical inference paths as would a human expert making the same prediction; its logic differs from doctrinal reasoning.

There are many benefits to using machine learning. While human input may be required at the outset (at least for supervised learning), the algorithm can be programmed to adapt over time without human intervention. Machine learning can manage larger datasets than may fall within the professional experience of any expert developing an expert system. Machine learning is thus particularly useful where datasets are large (settlements, sentencing for common crimes, patterns of criminal offending, student essays). Where datasets are small, machine learning tools are likely to overfit the data, meaning that they will 'learn' things that may be particular to a small sample of cases, thus misclassifying future cases or making unreliable predictions.

Machine learning has other limitations. It is only as useful as the data from which it draws; or, in other words, garbage in, garbage out. Quality here refers not only to the assumption that the underlying data are accurate, but also that they are representative, that there are no biases in the way they are collected, chosen and 'cleaned', and that there are no biases in how they are classified. In practice, these conditions are rarely met. Those using data-driven analysis are often constrained by the data that are available. Shortcuts include using social media (particularly Twitter) to gauge 'community' sentiment or using police crime data as a proxy for crimes actually occurring. Further, the learning process necessarily introduces inductive bias whenever the computer learner is asked to classify unseen examples.[13] The choice of algorithm affects the kind of bias, for example whether to prefer simpler or more complex inferences, false positives or false negatives, or which variables are used. When used for predictive modelling, machine learning assumes that past data can be used to predict the future, an assumption that can be disrupted by changes in law, policy, economic circumstances and other contextual elements. Because machine learning detects

[13] Tom M Mitchell, *Machine Learning* (McGraw-Hill, 1997) 39–45.

historical patterns and correlations, not causal relationships, the impact of an intervention or change in the system is hard to predict.

Perhaps the most significant danger of relying on data analytics is that correlations are misunderstood. A correlation between a subpopulation and high rates of offending may be spurious,[14] or can be caused by inherent characteristics of the subpopulation, by high rates of crime detection relating to that subpopulation, by different characterisations of activities of that subpopulation by police officers, or by historic treatment of that subpopulation (among other possibilities). The interventions one might want to propose ought to depend on one's belief about causation. However, data analytic tools are often used in a way that implicitly assumes a particular cause. For example, data may be used to assign a 'risk assessment' score to a criminal defendant. The algorithm may learn to give high scores to those within the subpopulation, given this pattern is evident in the historic data. If this is used to make decisions about individuals within that subpopulation, for example in relation to bail, sentencing and parole, then one is implicitly assuming a causal relationship between membership of the subpopulation and likelihood of reoffending that may be redundant.

Artificial intelligence in law thus comprises at least two categories of technologies. Expert systems are pre-programmed to mirror the questions and responses that might be given by a lawyer or law lecturer whereas machine learning identifies patterns in sufficiently large data sets to deduce rules (which need not look like legal rules) that can be applied to future examples. The former produces an output that looks like basic legal advice or a legal document, whereas the latter offers a classification or prediction without familiar explanatory reasoning.

[14] Tyler Vigen, *Spurious Correlations* (Hyperion, 2015).

3. Legal applications of artificial intelligence: What effect have they had?

A. Expert systems

In the 1970s and 1980s, there was significant interest in the development of computer-based tools for legal decision-making. The limits of these systems were evident to legal theorists familiar with issues of contradictory, circular, ambiguous, vague and contestable legal rules that often rely on social context and human interpretation.[15] While legal expert systems can provide basic legal advice where rules depend on simple criteria, they cannot provide advice in the same range of contexts as human lawyers. That does not mean they are useless – there are a variety of circumstances in which expert systems can answer important legal questions in high volumes (including tax calculations for basic PAYG earners, information on visa categories for which an individual may be eligible, and advice on federal benefit entitlements).[16] Where the factual matrix is finite and predictable, simple legal questions can be answered by pre-programmed logical steps. Expert systems are also an

[15] Philip Leith, 'The Rise and Fall of the Legal Expert System' (2010) 1(1) *European Journal of Law and Technology* <http://ejlt.org//article/view/14/1>. See generally Julius Stone, *Precedent and Law: Dynamics of Common Law Growth* (Butterworths, 1985) 63–74; Jeremy Waldron, 'Vagueness in Law and Language: Some Philosophical Issues' (1994) 82 *California Law Review* 509, 512–14; Laymen E Allen and Charles S Saxon, 'Analysis of the Logical Structure of Legal Rules by a Modernised and Formalised Version of Hohfeld Fundamental Legal Conceptions' in Antonio A Martino and Fiorenza Socci Natali (eds), *Automated Analysis of Legal Texts: Logic, Informatics, Law* (Elsevier, 1986) 385; Geoffrey Samuel, 'English Private Law: Old and New Thinking in the Taxonomy Debate' (2004) 24 *Oxford Journal of Legal Studies* 335, 362. See generally Ronald Stamper, 'Expert Systems – Lawyers Beware!' in Stuart S Nagel (ed), *Law, Decision-Making, and Microcomputers: Cross-National Perspectives* (Quorum Books, 1991) 19, 20.

[16] Eg See, eg, T J M Bench-Capon et al, 'Logic Programming for Large Scale Applications in Law: A Formalisation of Supplementary Benefit Legislation' in Thorne McCarty et al (eds), *Proceedings of the First International Conference on Artificial Intelligence and Law* (ACM Press, 1987) 190.

important component in many systems such as online legal databases and visualisation tools that support, rather than replace, the expertise of judges and practitioners.[17]

Renewed attention to legal expert systems can be seen in the increasing interest from law schools and law students. The University of Melbourne now offers a 'Legal Apps' course that teaches students how to develop a legal expert system using Neota Logic software. Through this course, students design a legal 'app' (essentially an expert system) for use in the not for profit sector.[18]

Experts systems can also be deployed in law schools, particularly in the context of online courses. The expertise of the lecturer is encapsulated in readings, videos and multi-media presentations, after which students are assessed against pre-programmed multiple choice or short answer questions. Expert systems can be used to build in progression paths, including alternative paths through the content based on responses to earlier questions. Where essays are written, they are generally still assessed by human academics.

The popularity of expert systems in law has thus increased slowly using technological tools that have remained largely constant, albeit with more modern interfaces and access through webpages and mobile applications (or 'apps').

---

[17] Graham Greenleaf, 'Legal Expert Systems: Robot Lawyers? An Introduction to Knowledge-Based Applications to Law' (Paper presented at the Australian Legal Convention, Sydney, August 1989), available at http://austlii.edu.au/cal/papers/robots89/; Richard Susskind, *The End of Lawyers? Rethinking the Nature of Legal Services* (Oxford University Press, 2008) 16; Graham Greenleaf, *Expert Systems Publications (The DataLex Project)* (30 December 2011), available at http://www2.austlii.edu.au/~graham/expert_systems.html; Floris J Bex et al, 'A Hybrid Formal Theory of Arguments, Stories and Criminal Evidence' (2010) 18 *Artificial Intelligence and the Law* 123, 125.

[18] The system used by Not-for-profit Law Information Hub <http://www.nfplaw.org.au/gettingstarted> was developed by students at the University of Melbourne.

B. Machine learning

One of the first areas of deployment of machine learning techniques in legal practice was in electronic discovery.[19] In particular, machine learning can be used to group documents together (such as near-duplicates), rethread email conversations, or identify documents that may be discoverable (on the basis of a training set classified by a lawyer). The latter function is usually achieved through a staged process, for example agreement on scope and protocol, human decisions on the discoverability of documents in a training set, machine learning and classification of remaining documents, quality assurance through random testing, possible retraining based on errors identified and reclassification, and possible repeat iterations through these latter steps. Ultimately, the machine learns to classify 'discoverability' based on features common among discoverable documents in the training set. Courts in the United States and United Kingdom have looked favourably on the use of machine learning technology in discovery to save costs in litigation and enhance consistency in the review process.[20] Actual performance of these tools has been variable,[21] and some criticise courts' overly enthusiastic embrace of these technologies.[22] Natural language processing and machine learning have also been used to extract relevant information from legal documents to

[19] For example, Precision Discovery <https://precisiondiscovery.com> and Deloitte's Dynamic Review <http://www2.deloitte.com/us/en/pages/advisory/articles/dynamic-review.html>.

[20] *Da Silva Moore v Publicis Groupe*, 287 FRD 182 (SD NY, 2012), *Rio Tinto Plc v Vale SA*, 306 FRD 125 (SD NY, 2015); *Pyrrho Investments Ltd v MWB Property Ltd* [2016] EWHC 256 (Ch). While the federal court is positively disposed towards the use of technology in discovery generally (see Federal Court of Australia, Practice Note CM 6), the issue of machine learning has not arisen specifically.

[21] Nicholas M Pace and Laura Zakaras, *Where the Money Goes: Understanding Litigant Expenditures for Producing Electronic Discovery* (Monograph, RAND Institute for Civil Justice, 2012) 62–6. <http://www.rand.org/content/dam/rand/pubs/monographs/2012/RAND_MG1208.pdf>

[22] Tonia Hap Murphy, 'Mandating Use of Predictive Coding in Electronic Discovery: An Ill-Advised Judicial Intrusion' (2013) 50 *American Business Law Journal* 609.

assist due diligence. As in the case of e-discovery, classification algorithms can identify important documents after suitable training.[23]

An interesting application of machine learning is the emerging field of litigation prediction. For example, the company Lex Machina makes a tantalising promise:

> We mine litigation data, revealing insights never before available about judges, lawyers, parties, and the subjects of the cases themselves, culled from millions of pages of litigation information. … With Lex Machina, you can easily view: … How likely is a judge to find infringement of a patent, fair use of a trademark or a Securities Act violation?[24]

The approach is promising; even a simple classification tree machine learning approach managed to beat experienced lawyers and scholars in its ability to predict the decisions of the United States Supreme Court (75% to 59%),[25] and the volume of data relied on in Lex Machina is significantly larger. While these developments are relatively recent, a neural network was used to estimate the quantum of damages for whiplash injury cases as early as the mid 1990s.[26]

---

[23] For example, eBrevia's *Diligence Accelerator* <http://ebrevia.com/diligence-accelerator/> and Deloitte's *Dynamic Review* <http://www2.deloitte.com/us/en/pages/advisory/articles/dynamic-review.html>.

[24] Lex Machina <https://lexmachina.com/what-we-do; https://lexmachina.com/legal-analytics>.

[25] Theodore W Ruger et al, 'The Supreme Court Forecasting Project: Legal and Political Science Approaches to Predicting Supreme Court Decision-Making' (2004) 104 *Columbia Law Review* 1150.

[26] Andrew Terrett, 'Neural Networks: Towards Predictive Law Machines' (1995) 3 *International Journal of Law and Information Technology* 94.

The idea of law as the ability to predict legal outcomes has a long history. Oliver Wendell Holmes famously defined law as '[t]he prophesies of what the courts will do in fact'.[27] This makes data analytics a more promising method of achieving artificial intelligence for legal realists than expert systems. As Loevinger wrote in the late 1940s, 'we have no terms to put into the machines' (or expert systems) because the significance of the 'vague verbalisations' used in legal rules is only 'ritualistic.'[28] If the law is what judges do, then legal prediction must be based on data. For followers of this line of thinking, Lex Machina knows the law better than the 'experts' as it understands, empirically, what judges actually *do*.

The promise of data analytics within the justice system is not limited to predicting the outcomes of civil litigation. Data analytics is also being used in some US jurisdictions to make decisions about bail, parole and sentencing. Using data about other accused, it is possible to correlate particular features of a defendant or the circumstances of the case to the situations of those who have broken bail conditions in the past.[29] Virginia has used a points system for granting parole to child sex offenders which, relying on statistical correlation with repeat offending, allows for earlier release of those whose victims were exclusively female.[30] A recent decision of the Supreme Court of Wisconsin has given the green light to the use of

[27] Oliver Wendell Holmes, 'The Path of the Law' (1897) 10 *Harvard Law Review* 457, 460-61.

[28] Lee Loevinger, 'Jurimetrics: The Next Step Forward' (1949) 33 *Minnesota Law Review* 455, 471.

[29] See Laura and John Arnold Foundation, 'Developing a National Model for Pre-trial Risk Assessment' (November 2013) <http://arnoldfoundation.org/sites/default/files/pdf/LJAF-research-summary_PSACourt_4_1.pdf>.

[30] Bernard E Harcourt, *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age* (University of Chicago Press, 2007) 10-15.

data-driven ‘risk assessment’ scores in sentencing, albeit not as the sole factor and not in determining the severity of the sentence.[31]

There is currently little direct use of data analytics in legal education, although it has been used extensively in marketing and evaluation. There are, however, possibilities on the horizon, including the automatic grading of essays based on similar technology to that used in e-discovery, namely training an algorithm to score essays based human scoring of essays in a training set.[32]

The extent of future use of artificial intelligence in the legal profession, legal academia and courts is hard to predict. The United States is already using data analytics for a variety of purposes, including settlement negotiations and risk-based calculations to assist in making decisions about bail, sentence and parole. However, the lack of familiarity of the approach taken by these technologies within legal professional communities may decrease adoption, particularly for sensitive areas such as the criminal justice system.[33] Adoption will also depend on the effectiveness of particular tools in solving particular problems,[34] as well as a broad range of factors that typically affect the diffusion of new technologies.[35] However, based on current trends, it is likely that over the next ten years there will be increasing deployment of legal expert systems based on current technology as well as increasing use of

---

[31] *Wisconsin v Loomis*, 881 NW 2d 749 (Wis, 2016).

[32] Salvatore Valenti et al, ‘An Overview of Current Research on Automated Essay Grading’ (2003) 2 *Journal of Information Technology Education* 319.

[33] Lyria Bennett Moses and Janet Chan, ‘Using Big Data for Legal and Law Enforcement Decisions: Testing the New Tools’ (2014) 37(2) *University of New South Wales Law Journal* 643.

[34] Patrice Flichy, *Understanding Technological Innovation: A Socio-Technical Approach* (Edward Elgar, 2007) 11.

[35] Everett M Rogers, *Diffusion of Innovations* (Free Press, 5th ed, 2003) 36, 219–66.

data analytic tools based on rapidly evolving machine learning techniques. However, as Part 4 will argue, the most crucial questions for predicting future impact are whether these technologies can provide what legal practitioners, academics and judges currently provide for clients, law students and society at large.

4. Projecting into the future: Possible effects over the next ten years

A recent book, *The Future of the Professions*, by Richard and Daniel Susskind, argues that expertise (including legal expertise) will in the future no longer require a similar number of human professionals.[36] This prediction relies on a range of technologies, including artificial intelligence. It builds on earlier work by Richard Susskind that explored how legal services would move from the bespoke model through stages of standardization, systematization, packaging and commoditization.[37] *The Future of the Professions* is oriented around the question, 'how do we share expertise in society?' The answer given by the authors is that this will increasingly use machines, operating on their own or with non-specialist users, rather than human professionals.

The Future of the Professions deploys an interesting analogy to illustrate how the professional model of provision of expertise might be disrupted. Drilling company executives are asked what product they sell.[38] The answer is not drills, but rather the hole in the wall. The point, of course, is that one can get to an end point in a different way. For lawyers and legal academics, the book argues, the 'hole in the wall' is the provision of legal expertise and

[36] Richard Susskind and Daniel Susskind, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Oxford University Press, 2015).

[37] Richard Susskind, *The End of Lawyers? Rethinking the Nature of Legal Services* (Oxford University Press, 2008).

[38] Susskind and Susskind, above n 36, 37.

legal education respectively. Human professionals are not necessarily the most efficient mechanism for the delivery of these. While professionals may be required to deliver 'bespoke' legal services and to undertake tasks involving moral deliberation and moral responsibility,[39] the numbers required will be significantly lower than the number of lawyers (and academics) employed today.

The argument is initially quite seductive. However, it assumes that technologies, including artificial intelligence, can offer the same 'hole in the wall' as legal practitioners, academics and judges. However, that is not always the case when viewed from the perspective of clients, law students and society at large.

A. Impact on clients

At least currently, online tools are not offering, nor claiming to offer, 'legal advice'. In fact, the terms and conditions on the various websites explicitly deny that this is what is being provided, for example:

> **LegalZoom.com**: 'the legal information contained on the Site and Applications is not legal advice and is not guaranteed to be correct, complete or up-to-date.'[40]

> **Thomson Reuters Cleardocs**: 'You agree that (a) we cannot, and do not, give you legal … advice; … (h) you must consult with a lawyer … or other appropriately qualified professional adviser … for advice concerning the suitability of a product you order using our service.'[41]

---

[39] Ibid, 249.

[40] LegalZoom Terms of Use <https://www.legalzoom.com/legal/general-terms/terms-of-use>.

[41] Cleardocs Terms and Conditions, <https://www.cleardocs.com/terms-and-conditions.html>.

This is unlikely to be a case of waiting for new technology that has greater capacity, but rather a case of limiting responsibility given the inherent limitations of expert systems, in particular the challenges of pre-programming all potential variations. Data analytics may, *if* done well, predict the probability of legal outcomes but it cannot explain its conclusions or provide advice on how to respond. What clients receive artificial intelligence tools will not, therefore, be substantively the same as legal advice.

Lawyers also take responsibility for the advice they give, and may be liable where advice is negligently given. On the other hand, sites offering automated services tend to disclaim both responsibility and liability for advice given. For example:

> **LegalZoom.com**: 'TO THE FULLEST EXTENT PERMITTED BY LAW, LEGALZOOM EXPRESSLY DISCLAIMS ALL WARRANTIES OF ANY KIND, WHETHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT.'[42]

> **Thomson Reuters Cleardocs:** '9. You agree that you indemnify us in relation to any cost, loss, liability or damage that any of you, your client, or a third party suffers (a) because the product you order is not suitable for its intended purpose …; (b) because you fail to obtain formal advice …'[43]

Thus what a person receives when seeking legal advice online using artificial intelligence technologies is not the *same* as what one gets from a lawyer giving legal advice. For some clients, a rough guide or draft document for which no responsibility is taken may be

[42] LegalZoom Terms of Use, above n 40.

[43] Cleardocs Terms and Conditions, above n 41.

sufficient. However, the difference will likely matter to those who rely, to their detriment, on faulty or unsuitable legal 'expertise'. It is thus important that users of these websites understand the consequences of relying on the information provided, bearing in mind the tendency of internet users generally to ignore 'terms and conditions' found on a linked webpage.

B. Impact on for law students

Expertise, unlike information, is not something that can easily be passed on to a third party. Obtaining a law degree requires more than learning and understanding a series of legal rules. If the goal is to feed legal information to the minds of receptive law students, one could run a series of online lectures, assigned readings and machine-graded quizzes. Even individualised learning can be seemingly automated with systems directing students along particular paths based on responses to earlier questions.

What a law student learns at law school, however, goes well beyond knowing legal rules. In recognition of this, law schools have moved beyond the 'Priestley 11' requirements for legal education to threshold learning outcomes for law degrees such as the LLB (TLOs).[44] The entire Priestley 11 list now fits within part of TLO 1, in particular the expectation that students are familiar with 'the fundamental areas of legal knowledge, the Australian legal system, and underlying principles and concepts.' The remaining learning outcomes focus on accomplishments that it would be difficult to gain without human facilitation, such as critical analysis (TLO 3(c)), creative thinking (TLO 3(d)) and professional judgment (TLO 2(d)).

---

[44] For example, Australian Learning and Teaching Council, Learning and Teaching Academic Standards Project, Bachelor of Laws, Learning and Teaching Academic Standards Statement (December 2010), <http://www.cald.asn.au/assets/lists/Education/LLB%20TLOsKiftetalLTASStandardsStatement2010%20TLOs%20LLB[2].pdf>.

In order to understand the difficulty of automated teaching of complex skills, consider TLO 5(a), the ability to 'communicate in ways that are effective, appropriate and persuasive for legal and non-legal audiences.' Students learn this skill through practice and feedback that extends throughout the degree. Students may be assessed for participation in class discussions or the giving of oral presentations, as well as for written essays, memoranda and advice. The feedback will encompass a range of factors, including clarity of expression, structure of argument (for written work) and the accuracy and insightfulness of the ideas expressed. This empowers students to improve over the course of a degree program.

Machine learning algorithms programs can learn to grade written work such as essays, but moving to such systems will destroy the process by which students learn to write well. Even where machine learning can *grade* an essay, it can only do so for basic features such as breadth of vocabulary, word and sentence count, use of different parts of speech, punctuation features, spelling and grammar.[45] These features may help distinguish among primary school students and may *correlate* with good writing more generally, but they are not the only features relevant to winning a competition such as this one. The gap between correlation and causation can be seen by imagining the consequences of disclosing the grading criteria. If one is aware of the features being assessed, one can write in a way that maximises one's score rather than addressing other desirables (such as structure, aesthetics or quality of reasoning). If one accepts correlates of quality as the basis for judging quality, one will find that savvy students will focus on one to the exclusion of the other, thus severing the original correlation.

This point about essay writing extends out to almost all of the TLOs. Only basic skills and core information may be learnt by watching videos and answering multiple choice questions with pre-programmed grading. Where expert systems are used in education, they

[45] See, for example, Valenti et al, above n 32.

are limited by the foresight of their designers. While machine learning avoids such human limitations, by relying on patterns and correlations in historic data, it ultimately teaches students to target correlates for quality, rather than quality itself.

C. Broader social impact

In their book, the Susskinds suggest that accepting unease about many new technologies expressed by professional bodies and professionals themselves are equivalent to leaving the rabbit to guard the lettuce.[46] Self-interest is said to bias professionals in favour of a stance that treats their own work as essential and irreplaceable. However, in the case of artificial intelligence in legal practice and legal academia, there are very real broader concerns.

Consider the relatively uncontroversial example of Lex Machina – the ability to predict the outcome of litigation in order to assist in developing a litigation strategy or deciding on an appropriate settlement offer. Such tools can have a positive impact on society, particularly to the extent that they encourage settlements. However, it also has an impact on the extent to which principles of law ultimately govern the settlement of disputes. Settlement currently takes place in the 'shadow of the law' – lawyers will consider their client's chances of success based on applying legal doctrine to the facts at hand.[47] Bargaining in the 'shadow of big data,'[48] relying on data-driven analytics and machine learning, has very different properties.[49] In particular, data can become skewed over time based on strategic decision-making, not legal principle. In civil litigation, an initial bias in favour of plaintiffs or

---

[46] Susskind and Susskind, above n 36, 32.

[47] Robert H Mnookin and Lewis Kornhauser, 'Bargaining in the Shadow of the Law: The Case of Divorce' (1979) 88 *Yale Law Journal* 950, 959–77.

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

[49] Bennett Moses and Chan, above n 33, 668.

defendants (or subcategories of either) in negotiating strength can be perpetuated through a belief, based on correlations in earlier settlement data, that damages are at a particular level. In criminal plea bargaining, historic biases against particular subpopulations can also be perpetuated through data-driven decision-making by prosecutors and defence lawyers.

However, these concerns about the impact on negotiations are less critical that concerns about the potential direct or indirect use of data-driven analytic tools by courts. As noted above, data analytics does not operate on the same logic as legal decision-making. The latter may not be perfect (if such were even possible for anyone other than Dworkin's Herculean judge),[50] but it rests on the idea that we are governed by rules and that decisions are made based on factors specified in statute or articulated through the process of developing the common law. Like cases are treated alike, with like-ness being determined according to the nature of the legal rule. Most of the factual landscape is irrelevant and inadmissible – my hair colour, my shoe size, and my ethnic origins are irrelevant (unless my shoe size is the same as that of a footprint at the scene of a crime). For machine learning, however, the approach is generally to use whatever works – prediction is the primary goal. As Berk and Bleich state in an article discussing the use of machine learning in parole decision-making:

> As a formal matter, one does not have to understand the future to forecast it with useful accuracy. Accurate forecasting requires that the future be substantially like the past. If this holds, and one has an accurate description of the past, then one has an accurate forecast of the future. That description does not have to explain why the future takes a particular form and certainly does not require a causal interpretation.[51]

[50] Ronald Dworkin, Law's Empire (Harvard University Press, 1986).

[51] Richard A Berk and Justin Bleich, 'Statistical Procedures for Forecasting Criminal Behaviour: A Comparative Assessment' (2013) 12 *Criminology & Public Policy* 513, 516.

They go on to suggest that if shoe size turns out to be a good predictor of recidivism, there is no reason to exclude it.

However, in criminal justice, predictive accuracy is not enough for two reasons. The first is that there may be causal explanations for this hypothetical correlation that render decisions spurious (for example, if those with large feet tend to be released earlier from prison historically due to the lack of available footwear). But the second reason is more crucial – legal rules specify factors to be taken into account in particular contexts because a socially legitimate institution (parliament or courts) have considered the issue and deemed those factors relevant. A failure to include shoe size is not a question of ignorance of a (hypothetical) correlation, but a conscious decision that the length of imprisonment ought not to be based on physical characteristics that a person cannot control and that are not of themselves dangerous.

In the United States, 'risk assessment' scores generated through the analysis of large data sets, are used in some jurisdictions to predict the future risk of re-offending, violent behaviour and absconding for bail applicants, parole decisions and, more recently, decisions as to how a sentence will be served.[52] The appropriateness of doing so in the context of sentencing was considered in the case of *Wisconsin v Loomis*.[53] The trial judge had referred to the risk assessment score of the defendant in the context of sentencing. Specifically, the circuit court stated

---

[52] Ariz Code of Judicial Admin § 6-201.01(J)(3) (2016); Idaho Code § 19-2517 (2016); Ky Rev Stat Ann § 532.007(3)(a) (2016); La Stat Ann § 15:326A (2016); Ohio Rev Code Ann § 5120.114(A)(1)-(3) (2015-2016); Okla Stat tit 22 § 988.18(B) (2016); Pa Cons Stat § 2154.7(a) (2016); Wash Rev Code § 9.94A.500(1) (2016).

[53] *Wisconsin v Loomis*, 881 NW 2d 749 (Wis, 2016).

> You're identified, through the COMPAS assessment [designed by Northpointe, Inc.], as an individual who is at high risk to the community. In terms of weighing the various factors, I'm ruling out probation because of the seriousness of the crime and because your history, your history on supervision, and the risk assessment tools that have been utilized, suggest that you're extremely high risk to reoffend.

The Supreme Court concluded that because the circuit court did not rely on solely on the risk assessment score and only used it in its determination regarding probation and not the severity of the overall sentence, the defendant's right to due process was not violated.[54] In particular, the court noted that while Loomis could not challenge the process through which the score was reached (given that it was a trade secret of Northpointe, Inc.), the court was provided with appropriate warnings about limitations, and Loomis was given an opportunity to verify some inputs as well as challenge the overall score by arguing that other relevant information should be taken into account.[55] There was an interesting discussion in the judgment about how the defendant's gender affected the outcome (in the context of the defendant's due process right not to be sentenced on the basis of gender). On that point, the court agreed with the State's assertion that inclusion of gender as a variable promotes accuracy given 'men and women have different rates of recidivism and different rehabilitation potential' so that 'a gender neutral risk assessment would provide inaccurate results for both men and women'.[56]

[54] Ibid at [8]-[9], [93], [98], [109]. But see [129] (concurring judgment of Roggensack CJ, seemingly implying that risk assessment tools can be considered in determining 'the sentence imposed' provided that the court did not rely on them).

[55] Ibid at [55], [56], [66].

[56] Ibid at [77], [86].

The controversy over the treatment of gender opens up questions about how other sensitive criteria, such as race, ought to be treated. The general tendency is to argue that race should not be used as a variable. However, this does not solve the problem and may, as the *Loomis* court suggested would be the case for ignoring gender, reduce accuracy. Further, even if race is ignored as a variable, there will be many other variables that may correlate with race, including socio-economic variables, education levels, where one lives and nature of family unit, even criminal history. One cannot remove all such variables since it would leave insufficient variables from which correlations could be drawn.

Race is a particular problem for 'risk assessment' algorithms in the United States. A ProPublica investigation into machine bias revealed the extent of differential treatment.[57] According to their analysis, black defendants were wrongly labelled as future criminals (meaning they were given a high risk assessment score but did not go on to commit future crimes) at almost twice the rate as white defendants. Even accounting for other factors such as criminal history and recidivism, age and gender, black defendants were more likely to be assessed as high risk than white defendants.[58] The differential impact can occur without using 'race' as a variable, indeed ignoring can in some circumstances make matters worse. For example, if black defendants overall have poorer educational outcomes (in the United States, this can be explained by their school funding model which relies on local taxes), then poor education outcomes may correlate less with criminality in black communities compared with white communities.[59] But, given African Americans are a minority of the population, the correlation for the whole population will impact on how minority defendants are scored. It is

---

[57] Julia Angwin et al, 'Machine Bias' *ProPublica* (23 May 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

[58] Ibid.

[59] Ibid.

possible that many of the factors potentially taken into account in risk assessment scoring, such as parent criminal histories, drug use and political views (the COMPAS questionnaire asks defendants whether they agree with the statement '[a] hungry person has a right to steal') could have a similar disparate impact.[60] Because the algorithm is a trade secret of the company Northpointe, Inc., it is difficult to assess whether the algorithm could be adjusted to remove potentially discriminatory impacts.

The question for the criminal justice system here is ultimately what features a 'risk assessment' tool would need to have so that its use in different contexts (bail, parole, sentencing) is both useful and consistent with fundamental community values concerning equal treatment before the law. Clearly predictive accuracy is essential, but cannot be sufficient. Avoiding discriminatory impact is far more complex than a decision to include or exclude variables or to compare subpopulations only against their 'peers'. An employer could not use statistics that demonstrate average poor performance of a particular racial group in a particular role to impact its hiring policies.[61] Should this same requirement operate when we move out of the private sphere into the public sphere of the criminal justice system? Should a tool be tested for differential impact on different racial groups, and what level would society be prepared to accept? How should an algorithm optimise not only false negatives versus false positives, but also the social impact of rating particular groups (for example, poor people or minorities) as more dangerous than others?

These questions are challenging ones, and the goal here is not to answer them, only to point out that the social consequences of using risk assessment tools in the criminal justice

[60] Ibid.

[61] *Racial Discrimination Act 1975* (Cth) s 15. See *also International Covenant on Civil and Political Rights*, opened for signature 16 December 1966, 999 UNTS 171 (entered into force 23 March 1976) art 2(1).

system are difficult to measure and potentially problematic. Just as technology is neither good nor bad but never neutral,[62] algorithms can be designed to achieve diverse objectives, but are never objective. Whenever accuracy is less than 100%, it is humans who decide what kind of errors are acceptable.

Because they cannot be objective, algorithms need to be made accountable to legal and policy objectives.[63] This requires lawyers and computer scientists to work together in design and evaluation. The importance of judges understanding these tools was emphasised in the judgment of Abrahamson J in the *Loomis* case, albeit dissenting on this point.[64] While one cannot delve directly into a human mind, we require human decision-makers in the judiciary and often in the bureaucracy to provide reasons for decisions, either publicly or to those affected. If such humans can rely on a non-transparent tool, such as a risk assessment score, in reaching a decision, then the mechanism through which such scores are produced ought to be understood. There are also important questions as to the right of a defendant to challenge the tool used in determining how a sentence will be served.[65] The answer given by the court in *Loomis*, being that the defendant could challenge the accuracy of the inputs, is unsatisfactory given the different hidden choices that might be made in analysing the data.

[62] Melvin Kranzberg, 'Technology and History: "Kranzberg's Laws"' (1986) 27 *Technology and Culture* 544, 545.

[63] Joshua A Knoll et al, 'Accountable Algorithms', (2017) 165 *University of Pennsylvania Law Review* (forthcoming).

[64] *Wisconsin v Loomis*, 881 NW 2d 749, [133]-[151] (Wis, 2016).

[65] Kate Crawford and Jason Schultz, 'Big Data and Due Process: Towards a Framework to Redress Predictive Privacy Harms' (2014) 55 *Boston College Law Review* 93; Danielle Keats Citron, 'Technological Due Process' (2008) 85 *Washington University Law Review* 1249.

Transparency is not always a feasible way to ensure accountability.[66] Some machine learning algorithms such as random forests and neural networks rely on complexity to enhance predictive accuracy, and this in turn reduces comprehensibility.[67] Where the black box cannot be revealed due to the technical complexity and emergent properties of machine learning processes, then one can use a complex, black box model as the basis for improving a simpler, comprehensible model,[68] one can derive a simplified explanation of the complex model,[69] or one can evaluate the black box model by testing it against a range of inputs.[70] The third possibility, evaluation against a range of inputs, provides an opportunity for testing not only overall predictive accuracy but also differential impact on subpopulations.[71]

Artificial intelligence tools ought to be treated with caution unless they can achieve the same socially desirable properties as the human process they seek partially to displace. This is not possible with current technologies, and is unlikely to become possible within the next ten years. However, unlike the limitations of artificial intelligence from the perspective of clients and law students, broader negative impacts will not necessarily affect uptake. The profession, and particularly judges, will need to understand the operation and limitations of artificial intelligence technologies if we are to avoid their dangers.

---

[66] Knoll, above n 63.

[67] Patrick Hall, 'Predictive modelling: Striking a balance between accuracy and interpretability' *O'Reilly* (11 February 2016), <https://www.oreilly.com/ideas/predictive-modeling-striking-a-balance-between-accuracy-and-interpretability>.

[68] Ibid.

[69] David R Warner Jr, 'A Neural Network-Based Law Machine: The Problem of Legitimacy' (1993) 2 *Information & Communications Technology Law* 135, 141.

[70] For example, see Anupam Datta et al, 'Algorithmic Transparency via Quantitative Input Influence', *Proceedings of 37th IEEE Symposium on Security and Privacy* (2016).

[71] Ibid.

5. Conclusion: Harnessing the benefits and limiting the detriments

The advantage of new technologies is that they are relatively malleable. Over time, it becomes harder to alter the socio-technical structures that harden around them.[72] We are still at the stage where today's decisions on the use of artificial intelligence in legal practice, legal academia and courts will have a significant impact on how these tools are constructed and perceived.

However, the development of artificial intelligence tools should not be unduly constrained. Those who are required to follow the law ought to have inexpensive access to its provisions so that efficiencies in communicating information, producing documents, running and settling litigation or conducting due diligence should be able to be pursued. Detriments referred to by the Susskinds, such as the loss of jobs for junior lawyers, are merely the latest event in the challenge that automation has posed since well before the Myers Committee released its report on the issue in 1980.[73] Regulation designed to preserve legal jobs alone would be counter-productive. In any event, if lawyers are taught new skills, such as the ability to assist in designing such systems, jobs in the new technology-enabled legal services sector will replace, to some extent at least, those in the old.

But claims about protectionism are not an answer to legitimate concerns about some applications of artificial intelligence in legal practice, legal academia and most critically in courts themselves. Websites providing legal services should be required to provide a clear, prominent statement as to the difference between what they provide and professional advice.

---

[72] Langdon Winner, 'Do Artifacts Have Politics?' (1980) 109(1) *Daedalus* 121, 127–8.

[73] Rupert H Myers, *Report of the Committee of Inquiry into Technological Change in Australia* (Australian Government Publishing Service, 1980).

University management need to remain aware of the limited role that fully automated online learning can play in disciplines such as law.

Data driven decision-making in criminal justice may seem like a very modern risk assessment approach, offering to eliminate potentially biased subjective decision-making. However, it raises fundamental questions about the factors that are allowed to play a part in decisions with significant implications for individuals and indeed society more broadly. Certainly the use of non-transparent tools that have not been subjected to a full evaluation for predictive accuracy, differential treatment and social impact should be rejected by courts and parole boards.

Artificial intelligence is likely to continue to have a significant impact on law, but the legal profession needs to guide the process. Legal professionals will need to understand the limitations of these kinds of tools, and ideally be involved in their design. There is much to be learnt from the way the military guide how humans and machines work together.[74] In particular, the military recognise that humans and machines each have things that they do well and poorly; they can work together to achieve important objectives if the appropriate delegations are made. For example, humans are better at responding to unanticipated events, and will thus continue to play an important role in strategic responses to new intelligence. In the legal context, artificial intelligence can play a useful role within a confined domain. Expert systems can help lawyers and the general public navigate unfamiliar legal terrain, while data analytics and machine learning can extract useful information from large datasets. However, both types of tools have limitations which can only be managed if well understood

[74] Committee on Integrating Humans, Machines and Networks: A Global Review of Data-to-Decision Technologies, Board on Global Science and Technology, Policy and Global Affairs, National Research Council of the National Academies (US), Complex Operational Decision Making in Networked Systems of Humans and Machines: A Multidisciplinary Approach (2014).

by the humans using them. The most crucial step in harnessing the benefits and limiting the detriments of these advances is thus education for judges, practitioners, legal academics and law students that encourages critical thinking about the use of artificial intelligence in law.

Bibliography

1. Books

Ronald Dworkin, Law's Empire (Harvard University Press, 1986)

Patrice Flichy, *Understanding Technological Innovation: A Socio-Technical Approach* (Edward Elgar, 2007)

Bernard E Harcourt, *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age* (University of Chicago Press, 2007)

Lawrence Lessig, *Code and Other Laws of Cyberspace* (Basic Books, 1999)

Tom M Mitchell, *Machine Learning* (McGraw-Hill, 1997)

Nicholas M Pace and Laura Zakaras, *Where the Money Goes: Understanding Litigant Expenditures for Producing Electronic Discovery* (Monograph, RAND Institute for Civil Justice, 2012)

Everett M Rogers, *Diffusion of Innovations* (Free Press, 5th ed, 2003)

Julius Stone, *Precedent and Law: Dynamics of Common Law Growth* (Butterworths, 1985)

Richard Susskind, *The End of Lawyers? Rethinking the Nature of Legal Services* (Oxford University Press, 2008)

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

Alan Tyree, *Expert Systems in Law* (Prentice Hall, 1989)

Tyler Vigen, *Spurious Correlations* (Hyperion, 2015)

2. Book chapters

Laymen E Allen and Charles S Saxon, 'Analysis of the Logical Structure of Legal Rules by a Modernised and Formalised Version of Hohfeld Fundamental Legal Conceptions' in Antonio A Martino and Fiorenza Socci Natali (eds), *Automated Analysis of Legal Texts: Logic, Informatics, Law* (Elsevier, 1986) 385

T J M Bench-Capon et al, 'Logic Programming for Large Scale Applications in Law: A

Formalisation of Supplementary Benefit Legislation' in Thorne McCarty et al (eds), *Proceedings of the First International Conference on Artificial Intelligence and Law* (ACM Press, 1987) 190

Mireille Hildebrandt, ‘A Vision of Ambient Law’ in Roger Brownsword and Karen Yeung (eds), *Regulating Technologies: Legal Futures, Regulatory Frames and Technological Fixes* (Hart, 2008).

Ronald Stamper, ‘Expert Systems – Lawyers Beware!’ in Stuart S Nagel (ed), *Law, Decision-Making, and Microcomputers: Cross-National Perspectives* (Quorum Books, 1991) 19

3. Articles

Lyria Bennett Moses and Janet Chan, ‘Using Big Data for Legal and Law Enforcement Decisions: Testing the New Tools’ (2014) 37(2) *University of New South Wales Law Journal* 643

Richard A Berk and Justin Bleich, ‘Statistical Procedures for Forecasting Criminal Behaviour: A Comparative Assessment’ (2013) 12 *Criminology & Public Policy* 513

Floris J Bex et al, ‘A Hybrid Formal Theory of Arguments, Stories and Criminal Evidence’ (2010) 18 *Artificial Intelligence and the Law* 123

Kate Crawford and Jason Schultz, ‘Big Data and Due Process: Towards a Framework to Redress Predictive Privacy Harms’ (2014) 55 *Boston College Law Review* 93

Anupam Datta et al, ‘Algorithmic Transparency via Quantitative Input Influence’, *Proceedings of 37th IEEE Symposium on Security and Privacy* (2016)

Oliver Wendell Holmes, ‘The Path of the Law’ (1897) 10 *Harvard Law Review* 457

Danielle Keats Citron, ‘Technological Due Process’ (2008) 85 *Washington University Law Review* 1249

Melvin Kranzberg, ‘Technology and History: “Kranzberg’s Laws”’ (1986) 27 *Technology and Culture* 544

Sylvia Kriven, ‘New IT System to transform justice’ (2014) 36 *Bulletin (Law Society of SA)* 9

Philip Leith, ‘The Rise and Fall of the Legal Expert System’ (2010) 1(1) *European Journal of Law and Technology* http://ejlt.org//article/view/14/1

Lee Loevinger, ‘Jurimetrics: The Next Step Forward’ (1949) 33 *Minnesota Law Review* 455

Robert H Mnookin and Lewis Kornhauser, ‘Bargaining in the Shadow of the Law: The Case of Divorce’(1979) 88 *Yale Law Journal* 950

Tonia Hap Murphy, ‘Mandating Use of Predictive Coding in Electronic Discovery: An Ill-Advised Judicial Intrusion’ (2013) 50 *American Business Law Journal* 609

Theodore W Ruger et al, ‘The Supreme Court Forecasting Project: Legal and Political Science Approaches to Predicting Supreme Court Decision-Making’ (2004) 104 *Columbia Law Review* 1150

Geoffrey Samuel, ‘English Private Law: Old and New Thinking in the Taxonomy Debate’ (2004) 24 *Oxford Journal of Legal Studies* 335

Dru Stevenson and Nicholas J Wagoner, ‘Bargaining in the Shadow of Big Data’ (2015) 67 *Florida Law Review* 1337

Andrew Terrett, ‘Neural Networks: Towards Predictive Law Machines’ (1995) 3 *International Journal of Law and Information Technology* 94

Alan Turing, ‘Computing Machinery and Intelligence’ (1950) 59(236) *Mind* 433

Salvatore Valenti et al, ‘An Overview of Current Research on Automated Essay Grading’ (2003) 2 *Journal of Information Technology Education* 319

Jeremy Waldron, ‘Vagueness in Law and Language: Some Philosophical Issues’ (1994) 82 *California Law Review* 509

David R Warner Jr, ‘A Neural Network-Based Law Machine: The Problem of Legitimacy’ (1993) 2 *Information & Communications Technology Law* 135

Langdon Winner, ‘Do Artifacts Have Politics?’ (1980) 109(1) *Daedalus* 121

4. Cases

Da Silva Moore v Publicis Groupe, 287 FRD 182 (SD NY, 2012)

Pyrrho Investments Ltd v MWB Property Ltd [2016] EWHC 256 (Ch)

Rio Tinto Plc v Vale SA, 306 FRD 125 (SD NY, 2015)

Wisconsin v Loomis, 881 NW 2d 749 (Wis, 2016)

5. Legislation, Regulations, Rules, Treaties

International Covenant on Civil and Political Rights, opened for signature 16 December 1966, 999 UNTS 171 (entered into force 23 March 1976) art 2(1)

Racial Discrimination Act 1975 (Cth) s 15

Ariz Code of Judicial Admin § 6-201.01(J)(3) (2016)

Idaho Code § 19-2517 (2016)

Ky Rev Stat Ann § 532.007(3)(a) (2016)

La Stat Ann § 15:326A (2016)

Ohio Rev Code Ann § 5120.114(A)(1)-(3) (2015-2016)

Okla Stat tit 22 § 988.18(B) (2016)

Pa Cons Stat § 2154.7(a) (2016)

Wash Rev Code § 9.94A.500(1) (2016)

Federal Court of Australia, Practice Note CM 6

6. Media

Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, 'Machine Bias' *ProPublica* (23 May 2016) <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

Patrick Hall, 'Predictive modelling: Striking a balance between accuracy and interpretability' *O'Reilly* (11 February 2016), <https://www.oreilly.com/ideas/predictive-modeling-striking-a-balance-between-accuracy-and-interpretability>

Marianna Papadakis, 'LawPath pursues growth in online legal services', *Australian Financial Review* (9 May 2014).

7. Unpublished papers

Graham Greenleaf, 'Legal Expert Systems: Robot Lawyers? An Introduction to Knowledge-Based Applications to Law' (Paper presented at the Australian Legal Convention, Sydney, August 1989) <http://austlii.edu.au/cal/papers/robots89/>

Graham Greenleaf, *Expert Systems Publications (The DataLex Project)* (30 December 2011) Austlii <http://www2.austlii.edu.au/~graham/expert_systems.html>

Joshua A Knoll et al, 'Accountable Algorithms' (2017) 165 *University of Pennsylvania Law Review* (forthcoming)

8. Reports

Australian Learning & Teaching Council, Learning and Teaching Academic Standards Project, Bachelor of Laws, Learning and Teaching Academic Standards Statement (December 2010), <http://www.cald.asn.au/assets/lists/Education/LLB%20TLOsKiftetalLTASStandardsStatement2010%20TLOs%20LLB[2].pdf>

Committee on Integrating Humans, Machines and Networks: A Global Review of Data-to-Decision Technologies, Board on Global Science and Technology, Policy and Global Affairs, National Research Council of the National Academies (US), Complex Operational Decision

Making in Networked Systems of Humans and Machines: A Multidisciplinary Approach (2014)

Rupert H Myers, *Report of the Committee of Inquiry into Technological Change in Australia* (Australian Government Publishing Service, 1980)

9. Online material

Arizona State University Online Programs <https://asuonline.asu.edu>

Axiom <http://axiomlaw.com>

Crowd & Co <https://crowdandco.com.au>

Deloitte's *Dynamic Review* <http://www2.deloitte.com/us/en/pages/advisory/articles/dynamic-review.html>

eBrevia's *Diligence Accelerator* <http://ebrevia.com/diligence-accelerator/>

Innovating Justice <www.hiil.org>

Laura and John Arnold Foundation, 'Developing a National Model for Pre-trial Risk Assessment' (November 2013) <http://arnoldfoundation.org/sites/default/files/pdf/LJAF-research-summary_PSACourt_4_1.pdf>.

Law Central Online <http://www.lawcentral.com.au>

LawAnswers <http://www.lawanswers.com.au>

LawPath <https://lawpath.com.au>

Legal Vision <https://legalvision.com.au>

legalzoom <http://www.legalzoom.com>

Lex Machina <https://lexmachina.com>

Modria <http://modria.com>

Neota Logic platform <http://www.neotalogic.com.au>

Not-for-profit Law Information Hub <http://www.nfplaw.org.au>

Precision Discovery <https://precisiondiscovery.com>

Riverview Law <http://www.riverviewlaw.com>

Thomson Reuters' *Cleardocs* <https://www.cleardocs.com>