2018: A Legal Research Odyssey: Artificial Intelligence as Disruptor*

Jamie J. Baker**

Cognitive computing has the power to make legal research more efficient, but it does not eliminate the need to teach law students sound legal research process and strategy. Law librarians must also instruct on using artificial intelligence responsibly in the face of algorithmic transparency, the duty of technology competence, malpractice pitfalls, and the unauthorized practice of law.

Introduction

¶1 My fascination with worker automation started at age twelve. My classmates and I traveled four hours away from our rural northern Michigan town of 2500, a town that had not changed much since the late 1800s when the manufacturing stronghold, the East Jordan Iron Works, was established. Most of our fathers

---

worked at the iron works; most of our mothers worked for Dura Automotive, a rural assembly line making component parts for the "Big Three" in Detroit. For many of us, this was our first big trip away from home. We were taking a three-day field trip to see, among other things, the world-renowned Henry Ford Museum. There were many memorable moments from this trip. I remember seeing the chair in which Lincoln was assassinated, with its blood-soaked back. I saw Buckminster Fuller's Dymaxion House. And I saw the future of automation in the auto industry.

¶2 One of the museum's exhibits displayed the new robotic arm of the automotive assembly line. The docent leading our school tour touted this as "revolutionizing" the line. As we filed to the next exhibit, I remember the distinct pit that formed in my stomach. While that robotic arm symbolized a revolutionary step in manufacturing, it also symbolized a loss of work and wages for the many struggling families in my hometown. The robotic arm would be great for Ford's bottom line; it would be disastrous for my family's bottom line.

¶3 Sure enough, within five years, Dura Automotive left East Jordan and took its jobs with it. While not solely attributable to automation, it was no doubt part of the equation. As a result of this early life experience, I developed a near obsession with prognostications about automation's future impact on society, including my chosen profession: law.

¶4 The assembly line involves the type of routinized work that is prime for automation, but we're now starting to hear about the automation of knowledge work in fields like finance, medicine, and law. And much of what we're hearing is that in the immediate future, knowledge work will see automation advances similar to those already seen in the manufacturing sector.

¶5 While it is naïve to think that automation won't affect knowledge work at all, it is clear that computing capability is not ready to replace highly skilled professionals. If stakeholders start to believe the hype of the PR campaigns surrounding artificial intelligence (AI) and automation, various sectors may be subject to premature disruption—the notion that workers are displaced before the technology is truly ready to replace them. To avoid premature disruption, legal professionals must understand current computing capability and the associated pitfalls of blindly relying on technology.

¶6 This article provides context for current computing capability and ultimately warns against the use of AI in violation of certain legal ethical obligations. Paragraphs 7–39 summarize current advances in AI technology and describe how knowledge-based professions such as finance, medicine, and law are using these advances. Paragraphs 40–51 discuss natural language processing (NLP) and the notion of premature disruption. Paragraphs 52–83 hypothesize about how legal research is likely to use AI while noting the complexities involved in the legal research process. That section ends by briefly discussing the ethical issues at play and the need to use AI responsibly, noting that law librarians are in the best position to teach prospective lawyers about the benefits and risks associated with the use of algorithms in law.

AI Becomes a Reality

¶7 To understand how AI will be employed in legal research and the various ethical implications at play, it is important to understand the current state of AI, particularly systems like DeepQA technology, and how professions such as finance, medicine, and law are already using this technology.

The Current State of Artificial Intelligence

¶8 Although AI has steadily progressed since the 1950s,[1] most software-driven capabilities still depend "on work processes that can be reduced to numbers and handled as mathematical calculations."[2] However, we are rapidly approaching a time when computing power will move beyond the reduction to numbers to the ability to process "vast quantities of text-based knowledge, and . . . [prove] able to answer questions that on their face have nothing to do with math and with high levels of reliability."[3] With some of the newer AI technologies, we are just starting to see this capability.

¶9 The first real iteration of the current capability was showcased in 2011, when IBM Watson beat former *Jeopardy!* champions Brad Rutter and Ken Jennings.[4] At this point, "IBM productized deep learning and natural language interaction to form a level of artificial intelligence known as 'cognitive computing.'"[5] To perform against the former *Jeopardy!* champions, Watson was programmed with basic language rules.[6] Additionally, Watson "also possesses over 100 separate modules with their own unique algorithm[s], each of which individually [tries] to determine the correct answers to questions on the show."[7] Watson is also made up of "a separate layer of algorithms that balance the results suggested by the computing modules to find the right answer."[8] Ultimately, Watson "combine[s] structured data, unstructured data, natural languages, and data analysis that could learn from other systems without the need for a human programmer to create software for every scenario."[9]

¶10 The genius of Watson is that "Watson does not generate one definitive answer but instead generates several possible answers, each with its own probability of being right."[10] During the *Jeopardy!* game, "Watson attempted to answer a question only if the probability of the top-ranked answer reached a certain threshold."[11]

---

1. *See generally History of Artificial Intelligence,* Wikipedia, https://en.wikipedia.org/wiki/History_of_artificial_intelligence [https://perma.cc/8VHA-NS3B].

2. Ray Worthy Campbell, *The Digital Future of the Oldest Information Profession* 3 (Jan. 17, 2016), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2716972 [https://perma.cc/VUV8-UTWH].

3. *Id.*

4. Ed Sohn*, alt.legal: Can Computers Beat Humans at Law?*, Above the Law: alt.legal (Mar. 23, 2016, 4:02 PM), http://abovethelaw.com/2016/03/alt-legal-can-computers-beat-humans-at-law/ [https://perma.cc/PHE2-84JC].

5. *Id.*

6. John O. McGinnis & Steven Wasick, *Law's Algorithm*, 66 Fla. L. Rev. 991, 1014 (2014).

7. *Id.*

8. *Id.*

9. Howard Lee, *Paging Dr. Watson: IBM's Watson Supercomputer Now Being Used in Healthcare*, 85 J. AHIMA 44, 44–47 (May 2014), http://library.ahima.org/xpedio/groups/public/documents/ahima/bok1_050656.hcsp?dDocName=bok1_050656 [https://perma.cc/WT7B-SB9A].

10. McGinnis & Wasick, *supra* note 6, at 1014.

11. *Id.*

¶11 IBM Watson is powered by "DeepQA" technology. In *Watson: Beyond Jeopardy!*, Ferrucci et al. provide the following explanation of DeepQA:

> DeepQA is a software architecture . . . informed by extensive research in question answering systems. . . . DeepQA analyzes an input question to determine precisely what it is asking for and generates many possible *candidate answers* . . . . For each of these candidate answers, a *hypothesis* is formed . . . . DeepQA searches its content sources for evidence that supports or refutes each hypothesis. For each evidence–hypothesis pair, DeepQA applies hundreds of algorithms that dissect and analyze the evidence along different *dimensions of evidence* . . . . The final result of this process is a ranked list of candidate answers, each with a confidence score indicating the degree to which the answer is believed correct, along with links back to the evidence.[12]

DeepQA remains flexible while using natural-language processing (NLP) to search large amounts of data.[13] Historically, "the ability to continuously process a stream of unstructured information from . . . [the] environment is . . . [something] for which humans are uniquely adapted. The difference . . . is that in the realm of big data, computers are able to do this on a scale that, for a person, would be impossible."[14]

¶12 In the coming years, DeepQA will be applied to many different domains. We're already starting to see this with IBM Watson–powered systems in areas as varied as medicine and cooking. Ferrucci et al. provide an example of adapting DeepQA to medicine, one of the first areas to adopt DeepQA computing. The authors illustrate three adaptations: content, training, and functional.

> [1] Content for the medical domain ranges from textbooks, dictionaries, clinical guidelines, and research articles, to public information on the web. There is often a tradeoff between reliability and recency of information available from these content sources.
>
> . . .
>
> [2] By using training questions [with known correct answers], the machine-learning models in DeepQA can learn what weight to attach to them. Alternatively, the decision maker may choose to do so manually, adjusting the confidence in a hypothesis based on its sources.
>
> . . .
>
> [3] Functional adaptation: DeepQA defines a general set of processing steps needed in a hypothesis evidencing system . . . . Conceptually, this pipeline includes analyzing and interpreting a question, searching, generating candidate hypotheses, retrieving supporting evidence, and finally scoring and ranking answers.[15]

Effectively, DeepQA's "language and knowledge processing infrastructure must . . . combine statistical and heuristic techniques to assess its own knowledge and produce its best answer with an accurate confidence—a measure of the likelihood it is correct based on a self-assessment of its sources, interference methods and prior performance."[16]

---

12. David Ferrucci et al., *Watson: Beyond Jeopardy!,* Artificial Intelligence, June–July, 2013, at 93, 94.

13. Martin Ford, Rise of the Robots 87 (2015).

14. *Id.*

15. Ferrucci et al., *supra* note 12, at 100–01.

16. Roger Leuthy, *IBM Announces Solid-State Memory Breakthrough*, Storage CH Blog (Dec. 28, 2010), https://rogerluethy.wordpress.com/author/rogerluethy/page/452/?chocaid=397 [https://perma.cc/GQ26-RZAH].

¶13 This capability to generate hypotheses and rank answers is unique to cognitive computing and DeepQA technology. IBM Watson, for example, does this "by analyzing the question as input, then generat[ing] a set of features and hypotheses by looking across data it has consumed as content. The computer then seeks the best potential response to the question."[17]

¶14 IBM Watson "[u]s[es] hundreds of reasoning algorithms embedded within the system . . . [to do] a deep comparison of the language of the question itself as well as each of the candidate answers."[18] The system then produces a relevance score that measures its confidence in the candidate answer.[19]

¶15 It is these types of expert systems, like IBM Watson powered by DeepQA technology, that are pushing the use of AI in the professions forward.

Artificial Intelligence in the Professions

¶16 For quite some time, the library world has harbored fears that technology may begin to replace human staff.[20] Within the last few years, and with the advent of DeepQA technology, the discussion surrounding the "world without work" has gotten louder—not just for librarians but for the professional world as a whole.[21]

¶17 According to Martin Ford's *Rise of the Robots*, Richard Susskind and Daniel Susskind's *The Future of the Professions,* and countless articles on point, nearly all professions are being bombarded with the message that they are doomed in the face of the AI boom.[22] "Computers are getting dramatically better at performing specialized, routine, and predictable tasks, and it seems very likely that they will soon be able to outperform many of the people now employed to do these things."[23]

¶18 In the early years of automation, particularly in the manufacturing sector, many could easily see how automation would transform the industry. The work of an assembly line was a prime target for automation as robots programmed to do routinized tasks were well suited to perform the same work previously done by humans and for far less money.

¶19 While routinized, predictable tasks were thought susceptible to automation, it has come as a surprise that much less routinized and predictable tasks are also being overtaken by automation. "In late 2013, two Oxford academics released a paper claiming that 47 percent of current American jobs are at 'high risk' of being automated within the next 20 years."[24] Automation threatens some occupations more than others:

---

17. Lee, *supra* note 9, at 44–47.

18. *Id.*

19. *Id.*

20. *See, e.g.,* Desk Set (Twentieth-Century Fox 1957) (offering an example from the 1950s of librarians in fear of being replaced by computers).

21. *See, e.g.*, Ford, *supra* note 13; Richard Susskind & Daniel Susskind, The Future of the Professions (2016); Derek Thompson, *A World Without Work*, Atlantic (July/Aug. 2015), https://www.theatlantic.com/magazine/archive/2015/07/world-without-work/395294/ [https://perma.cc/B2JW-JR4N].

22. *See generally* Ford, *supra* note 13; Susskind & Susskind, *supra* note 21; Thompson, *supra* note 21.

23. Ford, *supra* note 13, at 73.

24. Nathaniel Popper, *The Robots Are Coming for Wall Street*, N.Y. Times Mag. (Feb. 25, 2016), https://www.nytimes.com/2016/02/28/magazine/the-robots-are-coming-for-wall-street.html.

> [A]ccount software . . . can analyze and sort legal documents, doing the work that even well-paid lawyers often spend hours on. Journalists face start-ups like Automated Insights, which is already writing up summaries of basketball games. Finance stood out in particular: Because of the degree to which the industry is built on processing information—the stuff of digitization—the research suggested that it has more jobs at high risk of automation than any skilled industry, about 54 percent.[25]

These examples preview the extent to which knowledge-based work is vulnerable to automation. One of the more surprising fields affected is journalism. At one time, writing seemed least likely to be automated because it requires bespoke actions like retrieving information from a variety of systems; performing an often in-depth analysis; and writing understandable, compelling prose. But even it has been automated.[26] In fact, at least one scholar predicts that within fifteen years, more than ninety percent of news articles will be written algorithmically.[27]

¶20 Briefly reviewing automation's current impact on the finance and medical sectors will provide insight into the implications of automation on the legal sector.

Finance: Kensho and Beyond

¶21 To understand automation's effect on finance, we need look no further than Kensho, created by Daniel Nadler. Kensho parses an enormous number of datasets to provide predictive analysis for investors.[28] For example, when the Bureau of Labor Statistics released a monthly employment report, Kensho "scraped the data from the bureau's website. *Within two minutes*, an automated Kensho analysis [provided] a brief overview, followed by 13 exhibits predicting the performance of investments based on their past response to similar employment reports."[29]

¶22 Another practical example of Kensho's capability was its analysis helping investors understand how to position their portfolios in response to the Syrian civil war. "In the old days, [human workers] could draw on their own knowledge of recent events and how markets responded, . . . [or they] might have called a research analyst . . . to run a more complete study . . . ."[30] Now, with Kensho, the human workers can

> simply click an icon and . . . pick from a series of drop-down menus that narrow the search to a specific time period and a specific set of investments . . . .
>
> . . . *The whole process had taken just a few minutes. Generating a similar query without automation* . . . "would have taken days, probably 40 man-hours, from people who were making an average of $350,000 to $500,000 a year."[31]

Kensho works quickly and accurately by using cognitive computing to

> constantly [tweak] and broaden[] . . . . search terms, all with little human intervention . . . . Kensho's search engine automatically categorizes events according to abstract features . . . . The software . . . . looks for new and unexpected relationships between events and asset prices, allowing it to recommend searching that a user might not have considered. For this feature . . . . Nadler . . . . hired one of the machine-learning whizzes who worked on Google's megacatalog of the world's libraries.[32]

---

25. *Id.*

26. Ford, *supra* note 13, at 85–86.

27. *Id.* at 84–85.

28. Popper, *supra* note 24, at 1–2.

29. *Id.* (emphasis added).

30. *Id.* at 5–6.

31. *Id.* at 6–7 (emphasis added).

32. *Id.* at 6.

Ultimately, Kensho is performing tailored analysis once solely performed by highly educated, highly paid analysts.[33] And Kensho is but one computer program being used on Wall Street. "Machines are now responsible for most of the activity on Wall Street."[34] Other algorithms are being used to review stocks by looking at earnings statements, news reports, and regulatory filings because they are faster and "[a] lot can happen in [the] time frames before humans can react."[35]

¶23 The use of algorithms on Wall Street means greater efficiency and larger profits. But a major downside to this automation is that with markets reacting so quickly to computer analysis, the markets are more susceptible to glitches than ever before.[36] Critics argue that the markets are more volatile, and trading rules are not fit to handle orders in milliseconds.[37]

Medicine: IBM Watson for Medicine

¶24 Not only is finance seeing an insurgence of AI affecting its once-human processes, but medicine is too. After IBM Watson proved itself by winning at *Jeopardy!*,[38] the "supercomputer has moved on to practical applications—including being 'taught' to understand the complexities of healthcare."[39] A variety of "pilot programs . . . have recently launched that use Watson to improve healthcare processes and treatment" with its ability to combine structured and unstructured data to create several possible diagnosis options.[40]

¶25 The following case study illustrates the structured and unstructured data created in healthcare:

> A doctor gets a visit from a patient who has diabetes. The doctor determines he needs to do a blood sugar A1C test, a blood draw, an EKG, a blood pressure check, a cholesterol test, and a physical exam . . . . First, the results of a blood sugar test with a meter are usually logged in a patient's diary and not as part of a database. Since it's on paper, it is free text data and thus considered unstructured data. The A1C is done and logged into another system . . . . The blood draw goes to the lab, where technicians will look for abnormalities . . . . Blood pressure is usually done and hand written in a chart, creating more unstructured data that is not in the electronic health record (EHR). EKG results are checked by a doctor, but again stored as unstructured data in the health record. Finally, the physical exam results are typically written down by a doctor . . . and not entered as structured data in the EHR.[41]

This typical doctor-patient interaction shows that much of the medical data is unstructured.[42] This fact, together with "[t]he amount of medical information . . . doubling every five years,"[43] results in doctors' getting lost in data when trying to treat patients.

---

33. *Id.* at 10.

34. Michelle Fleury, *How Artificial Intelligence Is Transforming the Financial Industry*, BBC News (Sept. 16, 2015), http://www.bbc.com/news/business-34264380 [https://perma.cc/AP9N-K97J].

35. *Id.*

36. *Id.*

37. *Id.*

38. *See* Jennings Brown, *Why Everyone is Hating on IBM Watson—Including the People Who Helped Make It*, Gizmodo (Aug. 10, 2017, 8:45 AM), https://gizmodo.com/why-everyone-is-hating-on-watson-including-the-people-w-1797510888 [https://perma.cc/YW62-KNEQ] (discussing how IBM Watson works).

39. Lee, *supra* note 9, at 44.

40. *Id.*

41. *Id.* at 45.

42. *Id.*

43. *Id.*

¶26 Until recently, most computer programs in the healthcare arena stored and retrieved structured data. The systems were not programmed to understand natural language or analyze abstract data in an unstructured form.[44] But IBM Watson has changed all that. One of the first adopters of IBM Watson for healthcare was the premier cancer center of Memorial Sloan-Kettering Cancer Center (MSKCC), which taught

> IBM Watson about their breast and lung cancer research . . . and create[d] a system that [would] allow MSKCC to use the best available data to treat their cancer patients. IBM Watson used its cognitive computing natural language and decision support system to find patterns in unstructured information, mine patient data, analyze structured data, and look for disease patterns that most closely approximate each individual's case.[45]

The renowned cancer center of MD Anderson was another early adopter of the Watson technology, dubbing it "MD Anderson's Oncology Expert Advisor (OEA)." "By understanding and analyzing data in a patient's profile as well as information published in medical literature, the OEA can then work with a doctor to create evidence-based treatment and management options that are unique to that patient."[46]

¶27 The systems at MSKCC and MD Anderson both use DeepQA, a key characteristic of which is its use of search and NLP techniques. In addition, DeepQA "also helps ensure that the evidence provided in support of a set of possible solutions is readable and consumable by human users because the content is typically created by other experts in natural language rather than by knowledge engineers in formal rules."[47]

¶28 The Watson system for healthcare is a diagnostic support tool that uses "a rich set of observations about a patient's medical condition . . . and generates a ranked list of diagnoses (*differential diagnosis*) with associated confidences based on searching and analyzing evidence from large volumes of content."[48] Such diagnostic systems can help physicians avoid missing important potential diagnoses. But they are currently not widely used because they are not integrated into the day-to-day operations of healthcare organizations.[49] When a patient sees many different healthcare workers, and the patient's resulting medical data is scattered across different computer systems in both structured and unstructured form, it makes it nearly impossible for one program to have a complete picture of the patient's health record.[50] In addition, the diagnostic systems are difficult to interact with and the resulting list of possible diagnoses too long with little reasoning behind the diagnostic suggestions.[51] The diagnostic systems also do not provide an action plan for the physician because they are unable to ask for missing information that would increase confidence in a particular diagnosis.[52] Last, the diagnostic

---

44. *Id.*
45. *Id.* at 46.
46. *Id.*
47. Ferrucci et al., *supra* note 12, at 98.
48. *Id.* at 95.
49. *Id.* at 97.
50. *Id.*
51. *Id.*
52. *Id.*

systems are difficult to keep up to date, so the diagnostic suggestions are not always based on the latest medical evidence.[53]

¶29 To overcome some of these challenges, early adopters of IBM Watson as a clinical-decision support system have insisted on transparency in the decision-making processes that lead to the various diagnostic hypotheses.[54] Watson "must be able to decompose the confidence in a hypothesis into its constituent dimensions of evidence and compare [the evidence] across multiple competing hypotheses so that practitioners can arrive at their own conclusions."[55] Even given these transparency demands, the doctors working with Watson for Oncology[56] are critical of being able to validate Watson's results.[57]

¶30 While IBM Watson has made grand claims regarding its ability to revolutionize healthcare, in 2018 we are still waiting for the revolution. "AI can have a tremendous impact on health care and many other industries. But the technology doesn't seem advanced enough to have a transformational impact just yet."[58] In fact, recently the collaboration with the MD Anderson Cancer Center fell apart with the criticism that IBM was overly optimistic about Watson's abilities.[59]

Law

¶31 Recent AI applications in finance and medicine have increased the focus on AI capabilities for law.[60] From a historical perspective, law has not changed much since the industrial revolution. "The tools of research and of expression changed, progressing from goose quill to typewriters to word processors residing in the cloud, but throughout it all the nature of lawyers' daily work changed less than perhaps any other profession."[61] But as law shifts into the digital revolution, we will see the practice of law change dramatically. Many of the changes brought about by the digital revolution are just starting to take shape, and many others are still hidden from view.[62]

¶32 To date, expert systems have been developed for use by attorneys working in bankruptcy, immigration, estate planning, food and drug safety, and securities matters.[63] "In various ways, these expert systems evaluate input against existing information (including legal rules and analyses), draw inferences, make conclusions, and recommendations, and provide the reasoning therefor."[64] In addition, "[e]xpert systems and machine learning algorithms are even being used . . . to

---

53. *Id.*

54. *Id.* at 99.

55. *Id.*

56. *See Oncology & Genomics*, IBM Watson Health, https://www.ibm.com/watson/health/oncology-and-genomics/oncology/ [https://perma.cc/U3CG-GRDX].

57. *See* Brown, *supra* note 38.

58. *Id.*

59. David H. Freedman, *A Reality Check for IBM's AI Ambitions*, MIT Tech. Rev. (June 27, 2017), https://www.technologyreview.com/s/607965/a-reality-check-for-ibms-ai-ambitions/ [https://perma.cc/UF9P-W867].

60. *See generally* Joanna Goodman, Robots in Law: How Artificial Intelligence Is Transforming Legal Services (2016).

61. Campbell, *supra* note 2, at 1.

62. *Id.*

63. Pamela S. Katz, *Expert Robot: Using Artificial Intelligence to Assist Judges in Admitting Scientific Expert Testimony*, 24 Alb. L.J. Sci. & Tech. 1, 32 (2014).

64. *Id.*

advise judges. . . . While these expert systems don't make the decisions for the judges, they provide consultative or advisory tools to save time and provide consistency to decisions."[65]

> In Brazil, judges use a computer program that is programmed with an algorithm to review past decisions and recommend results in matters involving traffic collisions. Statistical software has been available to judges for many years to assist in sentencing, giving them an idea of sentencing on similar convictions in the past. Judges can now use expert systems to do that and more, such as: evaluate the convict's record, their seriousness, and frequency, as well as a number of other factors to be considered in sentencing. Then, these systems can weigh the factors and provide judges with the reasoning for their decisions. While these expert systems don't make the decisions for the judges, they provide consultative or advisory tools to save time and provide consistency to decisions.[66]

¶33 Improved technology has helped to augment lawyers' work in other ways too:

> Document assembly systems . . . help lawyers [draft] documents more quickly. Online research tools . . . have adopted . . . elements of artificial intelligence [that track] which returned sources are most heavily used and giving those sources more prominence in future searches for the same terms. Other tools . . . help lawyers hone in [on] the foundational cases more quickly.[67]

¶34 These advances have been works in progress as research tools, in particular, get to practical implementation against good data on a large scale. The "Big 2," LexisNexis and Westlaw, "have applied natural language processing (NLP) techniques to legal research for 10-plus years . . . . After all, the core NLP algorithms were all published in academic journals long ago . . . ."[68] These systems are continuously refining processes as computing power allows for a transition from natural language processing to natural language *understanding*.[69]

¶35 There has been incremental progress toward practical implementation against good data on a large scale through various vendors, such as Ravel Law[70] and Lex Machina.[71] For example, after Lex Machina built a large set of intellectual property (IP) case data, it used the corresponding data mining and predictive analytics techniques to forecast outcomes of IP litigation.[72] "Recently, it has extended the range of data it is mining to include court dockets, enabling new forms of insight and prediction."[73]

¶36 The next major step toward practical implementation will likely occur with a system like ROSS Intelligence, powered by IBM Watson, for legal research.[74]

---

65. *Id.* at 33.

66. *Id.*

67. Campbell, *supra* note 2, at 4.

68. Michael Mills, *Artificial Intelligence in Law: The State of Play 2016*, Thomson Reuters Legal Exec. Inst. 3 (Mar. 24, 2016), https://www.neotalogic.com/wp-content/uploads/2016/04/Artificial -Intelligence-in-Law-The-State-of-Play-2016.pdf [https://perma.cc/BRK5-S96E].

69. *See infra* ¶¶ 40–51; Mills, *supra* note 68, at 6.

70. Ravel Law, http://ravellaw.com/ [https://perma.cc/9VLF-F7EL] (recently purchased by LexisNexis).

71. Lex Machina, https://lexmachina.com [https://perma.cc/GUH9-CWMH] (recently purchased by LexisNexis).

72. Mills, *supra* note 68, at 5.

73. *Id.*

74. ROSS Intelligence, http://www.rossintelligence.com/ [https://perma.cc/BT4C-RL7M].

"ROSS uses machine learning technology to fine tune its research methods. The legal robot is accessed via computer and billed as a subscription service."[75] Promotional material for ROSS states:

> With the support of Watson's cognitive computing and natural language processing capabilities, lawyers ask ROSS their research question in natural language, as they would a person, then ROSS reads through the law, gathers evidence, draws inferences and returns highly relevant, evidence-based candidate answers. ROSS also monitors the law around the clock to notify users of new court decisions that can affect a case. The program continually learns from the lawyers who use it to bring back better results each time.[76]

Part of ROSS's learning process involves allowing users to upvote and downvote excerpts based on the robot's interpretation of the question. "Every time it answers a question, ROSS asks for feedback on its performance. Over time . . . ROSS's answers become more representative of the answers you would have gotten from the human professionals themselves. This is one of the primary features of all Watson progeny."[77]

¶37 ROSS's cofounder, Andrew Arruda, touts ROSS as saving lawyers up to thirty percent of their time, which, not coincidentally, corresponds with the same percentage that surveys show new attorneys spend on legal research.[78] ROSS is just starting to gain traction, with the law firm Baker & Hostetler announcing it would be licensing ROSS Intelligence to use in its bankruptcy practice. Other law firm subscribers include Latham & Watkins and von Briesen & Roper.[79]

¶38 In January 2017, Blue Hill Research Group released a benchmark report financed by ROSS, Inc., titled *ROSS Intelligence and Artificial Intelligence in Legal Research*.[80] Accordingly, the research objective was "[t]o assess the impact of ROSS-assisted use cases in bankruptcy law research with respect to: Information Retrieval Quality, Usability and User Confidence, and Research Efficiency."[81] Blue Hill used a panel of sixteen legal researchers to benchmark ROSS-use cases with those involving Boolean and natural language search capabilities of other research platforms.[82] Ultimately, the benchmark report found that when users conduct searches by entering questions in plain language, "ROSS's cognitive computing and semantic analysis

75. Karen Turner, *Meet "Ross," the Newly Hired Legal Robot*, Wash. Post (May 16, 2016), https://www.washingtonpost.com/news/innovations/wp/2016/05/16/meet-ross-the-newly-hired-legal-robot/?utm_term=.227b92ab403d [https://perma.cc/KMU8-MPBC].

76. *ROSS Intelligence Announces Partnership with BakerHostetler*, PR Newswire (May 5, 2016), http://www.prnewswire.com/news-releases/ross-intelligence-announces-partnership-with-baker hostetler-300264039.html [https://perma.cc/S6TG-FVF9].

77. *IBM Watson Takes the Stand*, Atlantic, http://www.theatlantic.com/sponsored/ibm-transformation-of-business/watson-takes-the-stand/283/ [https://perma.cc/8B8X-J9GL].

78. *The Tech Start-Up Planning to Shake up the Legal World*, BBC News (May 17, 2016), http://www.bbc.com/news/business-36303705 [https://perma.cc/98PQ-QYPV]; *Andrew Arruda: Artificial Intelligence and the Law Conference at Vanderbilt Law School*, YouTube (May 6, 2016), https://www.youtube.com/watch?v=LF08X5_T3Oc#t=2540.653469484.

79. Stephanie Francis Ward, *Jimoh Ovbiagele: Putting AI in Law Practice*, A.B.A. J. (Sept. 1, 2016, 11:45 AM), http://www.abajournal.com/legalrebels/article/jimoh_ovbiagele_profile/ [https://perma.cc/5PAC-THH7].

80. David Houlihan, ROSS Intelligence and Artificial Intelligence in Legal Research, Blue Hill Research (Jan. 17, 2017), (http://bluehillresearch.com/ross-intelligence-and-artificial-intelligence-in-legal-research/).

81. *Id.* at 1.

82. *Id.*

capabilities permit the tool to understand the intent of the question asked and identify answers 'in context' within the searched authorities."[83]

¶39 Based on the results, Blue Hill found that "ROSS AI plus Wexis outperforms either Westlaw or LexisNexis alone."[84] One of the primary takeaways from the Blue Hill report includes the following:

> It should be noted that none of these findings indicate that AI-assisted legal research constitutes a dramatic transformation in the use of technology by legal organizations. Rather, the use cases and impact reviewed indicate that tools like ROSS Intelligence more closely represent a significant iteration in the continuing evolution of legal research tools that began with the launch of digital databases of authorities and have continued through developments in search technologies.[85]

From this report, it is clear that ROSS powered by IBM Watson is a form of "augmented intelligence" that, guided by human experts, may make attorneys more efficient in their work.[86] Even with this efficiency, though, it is not ready to save attorneys thirty percent of their time because it does not have the computing capability to perform the requisite legal research.[87]

Natural Language Processing and Premature Disruption

¶40 ROSS Intelligence is described on its website as "an AI lawyer that helps human lawyers research faster and focus on advising clients."[88] And, as previously mentioned, ROSS's cofounder has touted that ROSS will save lawyers up to thirty percent of their time, the same percentage that surveys show new attorneys spend on legal research.[89] The various stakeholders, from firm partners to law students, must understand the current capabilities *and limitations* of ROSS and other expert systems, contrasted with developers' irresponsible hype, given the notion of premature disruption.[90]

¶41 While ROSS may well be better at NLP than any legal research system before it, ROSS is still limited by current computing capabilities.[91] "Sometime in the future it may be easier to find information without the help of a human individual interaction but that . . . seems a long way off because a computer has to be able to interpret whatever someone is saying and infer in ways that would be very challenging."[92]

---

83. *Id.* at 2.

84. Robert Ambrogi, *ROSS AI Plus Wexis Outperforms Either Westlaw or LexisNexis Alone, Study Finds*, Law Sites (Jan. 17, 2017), http://www.lawsitesblog.com/2017/01/ross-artificial-intelligence-outperforms-westlaw-lexisnexis-study-finds.html [https://perma.cc/95KM-QCFM].

85. *Id.* (emphasis added).

86. Jean P. O'Grady, *Hand in Hand with IBM Watson*, AALL Spectrum, Sept.-Oct. 2015, at 19, 20–21.

87. *See infra* ¶¶ 40–51.
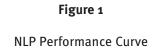
88. Ward, *supra* note 79.

89. *The Tech Start-Up Planning to Shake Up the Legal World, supra* note 78.

90. *See* Brian Sheppard, *Incomplete Innovation and the Premature Disruption of Legal Services*, 2015 Mich. St. L. Rev. 1797.

91. *See* Gary Marcus, *Artificial Intelligence Is Stuck. Here's How to Movie It Forward*, N.Y. Times (July 29, 2017), https://www.nytimes.com/2017/07/29/opinion/sunday/artificial-intelligence-is-stuck-heres-how-to-move-it-forward.html.

92. *The Tech Start-Up Planning to Shake Up the Legal World, supra* note 78.

Figure 1

NLP Performance Curve

¶42 When it comes to premature disruption, in medicine, IBM Watson has been criticized by Oren Etzioni as the "Donald Trump of the AI Industry" making outlandish claims about its ability that no credible data support.[93] Etzioni, CEO of the Allen Institute for AI, continued, stating that IBM Watson's "marketing and PR has run amok—to everyone's detriment."[94] This is because the technology is not truly ready to do what the PR folks tout that it can do. A Watson for Oncology designer opined that "IBM needs to be held accountable for the image that it's producing of its successes compared to what they're actually able to deliver, because at a certain point it becomes an ethical issue."[95]

¶43 The problem lies with the current limitations on computing capability, particularly with the ability of a computer to understand NLP. Looking at the NLP performance curve helps to explain this problem (see figure 1[96]). While NLP research has made great strides in producing artificially intelligent behaviors (e.g., Google, IBM's Watson, and Apple's Siri), none of such NLP frameworks actually understand what they are doing—making them no different from a parrot that learns to repeat words without any clear understanding of what the words mean.

---

93. *See* Brown, *supra* note 38.
94. *Id.*
95. *Id.*
96. Erik Cambria & Bebo White, *Jumping NLP Curves: A Review of Natural Language Processing Research*, IEEE Computational Intelligence Mag., May 2014, at 48, http://sentic.net/jumping-nlp-curves.pdf [https://perma.cc/GC2N-CD4W].

Today, even the most popular NLP technologies view text analysis as a word- or pattern-matching task. Trying to ascertain the meaning of a piece of text by processing it at word level, however, is no different from attempting to understand a picture by analyzing it at pixel level.[97] Thus far, NLP research has focused on word-level approaches. "Single-word expressions, however, are just a subset of concepts, multi-word expressions that carry specific *semantics and sentics.* Sentics . . . specifies the affective information associated with . . . real-world entities, which is key for common-sense reasoning and decision-making."[98] It is only with commonsense reasoning and decision making that NLP can truly leap from syntax to semantics and understand both "high- and low-level concepts as well as nuances in natural language understanding."[99] In practice, commonsense understanding allows a computer to properly deconstruct "natural language text into sentiments according to different contexts—for example, . . . the concept 'go read the book' as positive for a book review but negative for a movie review."[100]

¶44 As NLP systems continue to advance, they will gradually move from relying on syntactic, word-based techniques and start to exploit semantics more consistently and, hence, make a leap from the Syntactics Curve to the Semantics Curve.[101] But "[s]emantics . . . is just one layer up in the scale that separates NLP from true natural language *understanding*."[102] For systems to achieve the ability to accurately process information, the computational models must "be able to project semantics and sentics in time, compare them in a parallel and dynamic way, according to different contexts and with respect to different actors and their intentions."[103] This means that systems must eventually progress from the Semantics Curve to the Pragmatics Curve, "which will enable NLP to be more adaptive and, hence, open-domain, context-aware, and intent-driven."[104]

¶45 While the paradigm of the Syntactics Curve is the bag-of-words model and the Semantics Curve is characterized by a bag-of-concepts model, the paradigm of the Pragmatics Curve will be the bag-of-narratives model. In this last model, each piece of text will be represented by mini-stories or interconnected episodes, leading to a more detailed level of text comprehension and sensible computation. While the bag-of-concepts model helps to overcome problems such as word-sense disambiguation and semantic role labeling, the bag-of-narratives model will enable tackling NLP issues such as co-reference resolution and textual entailment.[105]

¶46 Pragmatics will provide a narrative understanding to allow for reasoning, decision making, and "sensemaking."[106] "Once NLP research can grasp semantics at a level comparable to human text processing, the jump to the Pragmatics Curve will be necessary, in the same way as semantic machine learning is now gradually evolving from lexical to compositional semantics."[107]

---

97. *Id.* at 51.
98. *Id.*
99. *Id.*
100. *Id.* at 51–52.
101. *Id.* at 56.
102. *Id.* at 52 (emphasis added).
103. *Id.*
104. *Id.*
105. *Id.*
106. *Id.* at 54.
107. *Id.* at 55.

¶47 Jumping the curve, however, is not an easy task. In fact, the origin of human language has been called the hardest problem of science.[108] "[A]lgorithms are limited by the fact that they can process only information that they can 'see.' Language, however, is a system where all terms are interdependent and where the value of one is the result of the simultaneous presence of the others."[109] As text processors, humans "'see more than what we see' in which every word activates a cascade of semantically-related concepts that enable the completion of complex NLP tasks, such as word-sense disambiguation, textual entailment, and semantic role labeling, in a quick and effortless way."[110] For an intelligent system to truly organize concepts into knowledge, the system must understand "physical knowledge of how objects behave, social knowledge of how people interact, sensory knowledge of how things look and taste, psychological knowledge about the way people think, and so on."[111] This type of computing power is still a long way off.[112]

¶48 As this detailed discussion of current computing capability suggests, "[w]hile exponential acceleration offer[s] valuable insight into the advance of information technology over a relatively long period, the short-term reality is more complex."[113] In the short term, progress will likely thrust forward but then stop "while new capabilities are assimilated into organizations and the foundation for the next period of rapid advance is established."[114]

¶49 In the long term,

> some . . . see computers continuing to double in power every two years, reaching levels of computing power by the 2020s that rival the human brain and that by the 2050s rival, in a single desktop machine, the power of all human brains combined. Even aside from the growth in processing power, there is every reason to expect that learning algorithms will wring ever-greater performance from existing machines. Given such vast increases in computational power, they see computers as besting humans at what lawyers do, which is to provide reliable, expert answers to difficult questions.[115]

¶50 In the short term, it is imperative that the various stakeholders consider the notion of premature disruption, whereby technologies replace human workers before the technology is truly ready to perform at the level of the replaced humans.[116] Even though some software is being touted as ready to replace humans,[117] in cases "where software must interact directly with people, . . . software has largely failed to leverage the advances that have occurred in hardware."[118] And DeepQA is just the latest software advancement to offer workplace assistance.[119]

---

108. *Id.* at 56.

109. *Id.*

110. *Id.*

111. *Id.*

112. *Id.* at 48.

113. Ford, *supra* note 13, at 64.

114. *Id.*

115. Campbell, *supra* note 2, at 8; *see also* Susskind & Susskind, *supra* note 21.

116. *See* Sheppard, *supra* note 90.

117. *See id.*

118. Ford, *supra* note 13, at 71; *see supra* ¶¶ 24–30.

119. J. Philip Craiger, *Technology, Organizations, and Work in the 20th Century*, Soc'y for Indus. & Organizational Psychol. (Jan. 1997), http://www.siop.org/tip/backissues/tipjan97/craiger .aspx [https://perma.cc/FP26-F5BC].

¶51 Experts generally agree that the greatest potential for immediate-future improvement is still in routine, repetitive tasks.[120] "For the next few decades . . . [it will be] a more complicated time—an interregnum in which the computers are not as smart as people but smart enough to do many of the tasks that make us money."[121] At this point, however, there is every reason to believe this process will continue to accelerate in the long term.[122]

AI in Legal Research

The algorithm can solve a case. It cannot build a case.[123]

¶52 The shorter-term reality means that legal research is nowhere near being automated. Legal research is not routine or repetitive. It is a highly sophisticated skill that requires a level of thinking better suited to the human brain. Christopher Columbus Langdell, dean of the Harvard Law School from 1870 to 1895, famously said that the law is a science, and the library is its laboratory.[124] From Langdell's time onward, law has seen significant improvements in the ability to access the vast trove of legal information.[125] It's literally at our fingertips. Currently, however, the process is at an interim period between giving lawyers access to information and truly providing relevant information, in a meaningful way, when it is sought.

DeepQA Applied to Legal Research

¶53 While lawyers can generally access the information they seek, computers do not yet have the ability to move beyond natural language processing to natural language *understanding*. It is impossible, then, for computers to truly perform effortless expert legal research. Expert legal research takes a level of creativity that requires context and pragmatic-level understanding to be performed properly.[126]

¶54 The skillful advocate strings together rules in a way that justifies the result she is seeking and at the same time encompasses the factual occurrence in a way that makes the rules she has selected appear to be the ones best applicable to the situation. The lawyer's research strategy is to identify the string of rules that both leads to a desired result and plausibly encompasses a set of facts that accounts for what has happened to her client.[127]

¶55 The advocate knows that "[t]he goal is not to reach the right decision but to make the best argument for one side."[128] And this requires a level of sophisticated pragmatic thinking that is distinctive from logic or scientific reasoning, for

---

120. Ford, *supra* note 13, at 71.

121. Popper, *supra* note 24, at 8.

122. *See* McGinnis & Wasick, *supra* note 6, at 1050.

123. Siddhartha Mukherjee, *A.I. versus M.D.,* New Yorker (Apr. 3, 2017), https://www.newyorker.com/magazine/2017/04/03/ai-versus-md [https://perma.cc/U7NA-BRJA].

124. *See* F. Allan Hanson, *From Key Numbers to Keywords: How Automation Has Transformed the Law*, 94 Law Libr. J. 563, 563, 2002 Law Libr. J. 563, ¶ 1.

125. *See supra* ¶¶ 31–39.

126. *See supra* ¶¶ 40–51.

127. Bruce G. Buchanan & Thomas E. Headrick, *Some Speculation About Artificial Intelligence and Legal Reasoning*, 23 Stan. L. Rev. 40, 47–48 (1970).

128. Phoebe C. Ellsworth, *Legal Reasoning*, *in* Cambridge Handbook of Thinking and Reasoning 685 (Keith J. Holyoak & Robert G. Morrison eds., 2005).

which AI is currently generally better suited.[129] As noted, to create the best argument for a client, the pragmatic-level, analytical thinking is inextricably linked to legal research.[130] The legal research process requires the highest level of NLP because it is "impossible to do legal research without analyzing, synthesizing, and applying the information found, both to the original issue and to the research plan developed to address the issue."[131] Legal research "cannot be mechanically divorced from legal analysis and reasoning."[132]

¶56 According to New York Law School's Kris Franklin, "[u]nderstanding how legal authorities are most effectively deployed to build legal arguments requires mastery of all of the most fundamental components of legal reasoning: reading sources of law meticulously, interpreting them critically, and applying them strategically."[133] Legal research, therefore, is linked directly to the "fundamental components of legal reasoning."[134] Moreover, "[i]f AI is to do justice to [legal research] processes, . . . it needs to accommodate their complexity in a realistic manner . . . . [L]aw, like many other areas, challenges AI to articulate the architectural features required to support reasoning in a domain saturated by complexity, uncertainty, defeasibility, and conflict."[135]

¶57 Because of the current limitations of NLP used in the existing retrieval systems, however, these systems help only with the periphery of this process. "They retrieve cases and statutes that are potentially relevant to some of the facts under consideration,"[136] but they do not produce the legal arguments that make up the end product of the research. That is because "legal search engines still work as a searchable index. Lawyers searching the index play a guessing game, trying to come up with the magical combination of terms that will get the search engine to return the relevant case law. The guessing game takes time, energy, and money."[137]

¶58 In the short term, as law adopts the use of DeepQA technology akin to the medical field's, the early iteration of the technology, taking into account current NLP capabilities, will allow "legal search engines to eliminate the guessing game by understanding, at a human level, the legal question being posed."[138] And "[i]nstead of typing in a search term [using Boolean connectors] . . . , the lawyer will simply ask 'find case law where the court discusses whether helping to cover up a conspiracy means you are responsible for the acts of the conspiracy.'"[139] The more sophisticated NLP capabilities using semantic-level understanding should be able to retrieve relevant results. Researchers need to understand, though, that relevant does not necessarily mean the "best" results to advocate for the client.

---

129. *Id.*

130. Sarah Valentine, *Legal Research as a Fundamental Skill: A Lifeboat for Students and Law Schools*, 39 U. Balt. L. Rev. 173, 209–10 (2010).

131. *Id.* at 210.

132. *Id.* at 211.

133. *Id.*

134. *Id.*

135. John A. Barnden & Donald M. Peterson, *Artificial Intelligence, Mindreading, and Reasoning in Law*, 22 Cardozo L. Rev. 1381, 1383 (2001).

136. Buchanan & Headrick, *supra* note 127, at 51.

137. McGinnis & Wasick, *supra* note 6, at 1018.

138. *Id.*

139. *Id.*

¶59 During the early stages, it is unlikely that a search engine will be able to determine, on its own, the one case that is most on point. "Instead, following Watson, the search engine will likely use competing algorithms to 'score' each possible case for how well it lines with the search query and come up with a short list of the top-ranked cases."[140] While "[t]he algorithm [will] . . . also take into account non-language related factors, such as whether the opinion was heavily cited to or searched for," it is, at this point, impossible for the legal search engines to choose the "best" case to make the most creative legal argument.[141]

The Limitations of AI and the Need to Use AI Responsibly

¶60 While DeepQA and its progeny have great potential to aid legal research, their current NLP capabilities limit their usefulness. This is a big problem for well-researched, well-reasoned legal analysis in a complex case.

> This [overall] problem [is in the] . . . limitations in software: (1) it cannot predict the infinite fact patterns that occur in difficult cases that are typically litigated; [and] (2) because machine learning is largely based on pattern recognition, it is likely to provide the easy solution associated with a similar easy case, while unable to replicate common sense judgments regarding important loopholes and policy concerns that apply to more specific fact patterns associated with substantially more difficult cases . . . . Th[e] limiting user interface, which would be the most likely artificial intelligence solution, could make the odds for the client even worse. For example, a client might [be advised to] settle based on a software program's [results] even though contract terms were ambiguous or unconscionable, not knowing that a court would not have upheld them [because those cases were not returned].[142]

While DeepQA is a good start in finding relevant cases, it is the harder "cases, ones that do not occur regularly and are generally not predictable, [that] clients now decide to consult with an attorney [about]. For easy cases, that occur regularly, like a standard rear-end collision with no personal injury, many people settle without attorneys."[143] For this reason, DeepQA and other AI agents are currently of limited use in legal research.

¶61 The technology will inevitably continue to evolve and advance at an exponential rate, and attorneys must understand the issues surrounding computing capability. There is a real danger in relying blindly on algorithms to do sophisticated legal research without understanding how the algorithms generate results. Lawyers must be cognizant of total reliance on algorithms in the face of algorithmic accountability, the Duty of Technology Competence, malpractice pitfalls, and the unauthorized practice of law.

Algorithmic Accountability and Computational Negligence

¶62 The danger with algorithmic accountability is that currently little regulation in this area exists. "As routine matter of business, the corporations offering legal services do not share the[ir] algorithms."[144] Without understanding how the

---

140. *Id.*

141. *Id.*

142. Frank Pasquale & Glyn Cashwell, *Four Futures of Legal Automation*, 63 UCLA L. Rev. Discourse 26, 43 (2015).

143. *Id.* at 44.

144. Campbell, *supra* note 2, at 11.

algorithms generate results, it is difficult, if not impossible, for attorneys to vet the information.

> [I]nvisible to the user, these products could be subject to intentional or unintentional biases. For example, a product relying on Big Data analysis and statistical correlation might give different advice in response to a criminal charge if race or income were a variable, embedding, unknown to the consumer, historic biases in the information given. With the algorithm hidden, the bias would be, as a practical matter, undetectable. In order for digital legal services to achieve their potential, these issues need careful thought. Given the lack of transparency, it cannot be assumed that the market will provide a sufficient check as consumers may not even realize the issue exists, and most likely will be unable to evaluate the options, even if they recognize the general issue of private algorithms. While tight government regulation of the giant Internet companies summons up its own parade of horribles, the issue is too important to ignore.[145]

¶63 In addition, as we continue to transition from the Digital Age to the Algorithmic Society, algorithms will increasingly be used to govern populations.[146] The underlying data and algorithms will be used to understand, analyze, control, direct, order, and shape society. "Because the relationship is one of governance, the obligations are fiduciary."[147] Without algorithmic accountability and transparency,[148] the very people who create laws to govern will not be able to act in their fiduciary capacities.

¶64 To ensure that lawyers are able to meet their fiduciary responsibilities, the three principles of the Algorithmic Society should be taken into account:

> (1) With respect to clients, customers, and end-users, algorithm users are *information fiduciaries*.
> (2) With respect to those who are not clients, customers, and end-users, algorithm users have *public duties*. If they are governments, this follows from their nature as governments. If they are private actors, their businesses are affected with a public interest, as constitutional lawyers would have said during the 1930s.
> (3) The central public duty of algorithm users is to avoid externalizing the costs (harms) of their operations. The best analogy for the harms of algorithmic decision-making is not intentional discrimination but socially unjustified pollution.[149]

With respect to principal (1), fiduciaries "who use robots, AI agents, and algorithms have duties of good faith and trust toward their end users and clients. Fiduciary duties apply whether a business or entity uses robots, AI agents, or machine learning algorithms in delivering services."[150] In adopting this fiduciary duty, it is recognized that "the use of algorithms can harm not only the end-user of a service, but many other people in society as well."[151] It equates to "the socially unjustified use of computational capacities."[152] Because "the algorithm doesn't have intentions, wants, or desires . . . . , we have to focus on the social effects of the use of a particular

---

145. *Id.*
146. Jack Balkin, *The Three Laws of Robotics in the Age of Big Data*, 78 Ohio St. L.J. 1217, 1226 (2017).
147. *Id.*
148. *See* Sara Castellanos & Steven Norton, *Inside Darpa's Push to Make Artificial Intelligence Explain Itself*, Wall Street J. (Aug. 10, 2017, 12:59 PM), https://blogs.wsj.com/cio/2017/08/10/inside-darpas-push-to-make-artificial-intelligence-explain-itself/.
149. Balkin, *supra* note 146, at 1227.
150. *Id.* at 1230.
151. *Id.* at 1232.
152. *Id.* at 1233.

algorithm, and whether the effects are reasonable and justified from the standpoint of society as a whole."[153] And this duty is nowhere greater than in law, where the use of algorithms has the greatest ability to result in deleterious effects on society. Without algorithmic transparency and the ability to monitor operations, there is no means to provide a rebuttal or method for holding the algorithm accountable. This will not do when so much is at stake.[154]

¶65 Even with sufficient transparency, biases may still exist in algorithmically organized systems.[155] This could be a result of "the algorithm creators . . . build[ing] into their creations their own perspectives and values."[156] Or it could be that "the datasets to which algorithms are applied have their own limits and deficiencies."[157] Realistically, "[t]he algorithms will be primarily designed by white and Asian men—the data selected by these same privileged actors—for the benefit of consumers like themselves."[158] Thus, "[t]he makers of these algorithms and the collectors of the data used to test and prime them have nowhere near a comprehensive understanding of culture, values, and diversity."[159] A prime example of this dangerous bias exists "in criminal justice, for example, [where] . . . an algorithm that fulfills basic statistical desiderata is also a lot more likely to rate black defendants as high-risk even when they will not go on to commit another crime."[160]

¶66 Another core problem with algorithmic-based decision-making and transparency is that the machines have literally become black boxes—even the developers and operators do not fully understand how outputs are produced. "There is a larger problem with the increase of algorithm-based outcomes beyond the risk of error or discrimination—the increasing opacity of decision making and the growing lack of human accountability."[161] Thus,

> [t]he danger in increased reliance on algorithms is that the decision making process becomes oracular: opaque yet unarguable. The solution is design. The process should not be a black box into which we feed data and out comes an answer, but a transparent process designed not just to produce a result, but to explain how it came up with that result. The systems should be able to produce clear, legible text and graphics that help the users—readers, editors, doctors, patients, loan applicants, voters, etc.—understand how the decision was made. The systems should be interactive, so that people can examine how changing data, assumptions, rules would change outcomes.[162]

¶67 In the simplest terms, the call for algorithmic transparency would allow sophisticated users to "review a software-driven action after-the-fact . . . to see if it

---

153. *Id.* at 1234.

154. *Id.* at 1239.

155. Janna Anderson & Lee Rainie, *Code-Dependent: Pros and Cons of the Algorithm Age*, Pew Res. Ctr. 11 (Feb. 8, 2017), http://www.pewinternet.org/2017/02/08/code-dependent-pros-and -cons-of-the-algorithm-age/ [https://perma.cc/P3QF-LW76].

156. *Id.*

157. *Id.*

158. *Id.* at 12.

159. *Id.* at 21.

160. Emma Pierson, *Hey Computer Scientists! Stop Hating on the Humanities*, Wired (Apr. 24, 2017, 10:00 AM), https://www.wired.com/2017/04/hey-computer-scientists-stop-hating-humanities/ [https://perma.cc/J3KU-2B6L].

161. Anderson & Rainie, *supra* note 155, at 19.

162. *Id.* at 22–23.

comports with applicable social, political, or legal norms."[163] This is particularly challenging for machine learning systems that adapt on their own and change often without the ability to create a discernable decision log.[164]

> Although the subject of such a process may not have the literal ability to know or understand what reasons are behind a decision, when a sensitive decision is being made—or . . . when the state is making decisions that raise due process concerns—the state must use software that furnishes relevant evidence to support evaluation and hence allow for technical accountability.[165]

¶68 Law, like medicine, should require algorithmic transparency. Arguably, all legal outcomes are sensitive decisions. If users do not have a clear picture of how a decision was made or how a particular case hypothesis was generated, the user cannot fulfill his or her fiduciary duty to the client and avoid "computational negligence."

The Duty of Technology Competence and Malpractice Pitfalls

¶69 Not only is there a theoretical fiduciary duty created by the use of algorithmic decision making when governing populations, there is also the very real, newly created Duty of Technology Competence.

> [T]he American Bar Association formally approved a change to the Model Rules of Professional Conduct to make clear that lawyers have a duty to be competent not only in the law and its practice, but also in technology. More specifically, the ABA's House of Delegates voted to amend Comment 8 to Model Rule 1.1, which pertains to competence, to read (emphasis added)
>
> > ". . . . To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, *including the benefits and risks associated with relevant technology*, engage in continuing study and education and comply with all continuing legal education requirements to which the lawyer is subject."
>
> This being a model rule, it must be adopted in a state for it to apply there . . . . So far, twenty-one states have done so.[166]

Because most attorneys do not have specialized training focused on a particular technological field, basic ethical rules provide a framework for determining a practitioner's professional duties and obligations with regard to technology—specifically, rules pertaining to competent client representation, adequate supervision, confidentiality, and communications.[167] Thus far, the Duty of Technology Competence has been interpreted to apply to eDiscovery.[168] While there has been

---

163. Deven R. Desai & Joshua A. Kroll, *Trust But Verify: A Guide to Algorithms and the Law*, 31 Harv. J.L. & Tech. 1, 39 (2017).

164. *Id.* at 49–50.

165. *Id.* at 44–45 (footnote omitted).

166. Robert Ambrogi, *Ethics and Technology Competence*, Above the Law: This Week in Tech (July 11, 2016, 3:02 PM), http://abovethelaw.com/2016/07/this-week-in-legal-tech-ethics-and -technology-competence/.

167. Stacey Blaustein, Melinda L. McLellan & James A. Sherer, *Digital Direction for the Analog Attorney-Data Protection, E-Discovery, and the Ethics of Technological Competence in Today's World of Tomorrow*, 22 Rich. J.L. & Tech. no. 4, 2016, 10, 25.

168. *See* State Bar of Cal. Standing Comm. on Professionalism, Responsibility & Conduct, Formal Op. No. 2015-193, https://www.calbar.ca.gov/Portals/0/documents/ethics/Opinions/CAL%20 2015-193%20%5B11-0004%5D%20(06-30-15)%20-%20FINAL.pdf [https://perma.cc/7ZUG-V7A9].

no guidance issued on the use of algorithms, it is not far-fetched to conceive that the use of algorithms will fall under this ethical rule at some point.[169]

¶70 In addition to the Duty of Technology Competence, ABA Model Rule 5.1 bears on a lawyer's duties regarding technology insofar as tasks aided or supported by technology are performed by someone other than the attorney. This responsibility extends to immediate as well as remote support staff, with ABA Model Rule 5.1 requiring that "[l]awyers must also supervise the work of others to ensure it is completed in a competent manner."[170]

> This attempt at establishing "the principle of supervisory responsibility without introducing a vicarious liability concept" has led to considerations regarding inexperience generally, but the implications for technological applications should be clear—an associate or other paralegal professional is much more likely to use technology to support legal work than she is to make a representation before a court or like body.[171]

¶71 Like ABA Model Rule 5.1, ABA Model Rule 5.3 sets forth responsibilities of partners and supervising attorneys to nonlawyer assistants. This rule

> further reinforces the responsibilities attorneys have to apply sufficient care in their practice when outsourcing supporting legal work to inexperienced non-professionals, and to ensure that confidentiality is maintained with outsourcing staff. This is not just a matter of supervising specific tasks. It also contemplates knowing which tasks are appropriate for delegation, both within the firm and to third-party vendors. For example, if a delegate of the attorney uses technology to begin an engagement, it's possible that such an arrangement could be viewed as "establish[ing] the attorney-client relationship," which may be prohibited under ABA Model Rule 5.5.[172]

¶72 On a practical level, lawyers could eventually use such technologies to replace lower-level legal professionals. For example, a software application could first conduct a fact-gathering intake session to formulate the questions that need to be answered. Then using algorithms that employ natural language understanding, the algorithm would analyze the user inputs to understand the question. The algorithms would generate the appropriate case law, statutes, and regulations to analyze and compile a memo that succinctly describes the current law.[173] It is not difficult to see that this would be considered the outsourcing of legal support work to a nonlawyer. And the lawyer is required to supervise the legal work accordingly.

¶73 The various ethical duties, including the Duty of Technology Competence, are inextricably linked to malpractice considerations. "The legal . . . industr[y is], by nature, [an] industr[y] of precision. A small typographical error in a legal document could result in a malpractice lawsuit . . . ."[174] Most tasks in this industry

---

169. *See generally* Lauren Kellerhouse, *Comment 8 of Rule 1.1: The Implications of Technological Competence on Investigation, Discovery, and Client Security*, 40 J. Legal Prof. 291 (2016).

170. Samantha Ettari & Noah Hertz-Bunz, *Ethical E-Discovery: What Every Lawyer Needs to Know*, Legal Tech. News (Nov. 10, 2015, 10:21 A.M.), https://www.law.com/legaltechnews/almID /1202742064964/?slreturn=20180003193746 (archived at Lexis Advance, https://advance.lexis.com /api/permalink/e04c5f11-08cc-4586-8119-9de6ee059246/?context=1000516).

171. Blaustein, McLellan & Sherer, *supra* note 167, at 28 (footnote omitted).

172. *Id.* at 28–29.

173. Justin D. Leonard, *Cyberlawyering and the Small Business: Software Makes Hard Law (But Good Sense)*, 7 J. Small & Emerging Bus. L. 323, 346 (2003).

174. *Can An A.I. Robot Perform Due Diligence Better Than Your Lawyer?*, Merger Technology.com (Sept. 28, 2016, 9:00 AM), http://mergertechnology.com/security/can-an-a-i -robot-perform-due-diligence-better-than-your-lawyer-3699 [https://perma.cc/5KZ8-82WB].

require meticulous attention to detail. Delegating the tasks to a computerized program involves a significant amount of trust.[175] Given the various ethical duties and the precision necessary to practice law, if the algorithms do not provide the requisite transparency, these duties, when violated, may open lawyers up to malpractice claims.

Unauthorized Practice of Law

¶74 Along with the issues that flow from algorithmic accountability, various legal ethical duties, and malpractice pitfalls, stakeholders must be cognizant of issues surrounding the unauthorized practice of law when using algorithms. "Ostensibly, the main policy rationale for this prohibition is legal services quality assurance and the protection of the public from unqualified legal practitioners, who while appropriating the full benefits of legal practice, often eschew the corresponding responsibilities that traditionally underpin the attorney and client relationship."[176]

¶75 To understand what constitutes the unauthorized practice of law, Texas has a preliminary definition that lists circumstances under which a person would be presumed to be practicing law.[177]

> These include giving advice or counsel to persons on their legal rights and responsibilities or to those of others, selecting, drafting, or completing legal documents or agreements that affect the legal rights of others, representing a person before an adjudicative body, including but not limited to documents preparation, or filing, or conducting discovery, or negotiating legal rights or responsibilities on behalf of a person.[178]

Anyone engaging in the unauthorized practice of law is subject to criminal and civil penalties.[179] After Texas adopted this preliminary definition, the ABA recommended that every state and territory adopt a similar definition.[180]

¶76 Algorithms, like websites, do not "just grow out of thin air and . . . aren't maintained out of thin air. They're put together by people . . . . It's the people who develop [the algorithms] that [arguably] provide the assistance."[181] The more the software, or algorithm, does, the greater the chances that it could be seen as conducting the unauthorized practice of law. For example, if "[t]he software . . . go[es] far beyond providing clerical services . . . . [to] determin[ing] where (particularly, in which schedule) to place information provided by the debtor, select[ing] exemptions for the debtor[,] and suppl[ying] relevant legal citations, ... [p]roviding such personalized guidance has been held to constitute the practice of law."[182]

¶77 Developers who create legal algorithms for nonlawyers must understand the distinction between clerical work and preparing legal documents, as this "is the traditional benchmark for ascertaining whether a non-attorney is engaged in unau-

---

175. *Id.*

176. Taiwo A. Oriola, *The Use of Legal Software by Non-lawyers and the Perils of Unauthorised Practice of Law Charges in the United States: A Review of the Jayson Reynoso Decision*, 18 Artificial Intelligence & L. 285, 286 (2010).

177. *Id.* at 307.

178. *Id.*

179. *Id.*

180. *Id.*

181. *Id.* at 291.

182. *Id.* at 295.

thorized practice of law."[183] The developers could be liable for the unauthorized practice of law if the legal software in question goes beyond simple clerical work to the drafting of legal documents or the proffering of legal advice.[184]

Algorithmic Literacy: Legal Research Instruction Implications

¶78 It behooves law librarians to bring these issues surrounding the use of algorithms to light during legal research instruction. In the words of Professor Robert Berring, "[i]n the midst of an information revolution that it cannot stop and seems hardly to understand, the legal profession must reassess the very way it thinks about legal research and legal research training."[185] "[N]atural language search, as it is refined, will have fundamental implications for legal search and ultimately the form of law."[186]

¶79 Using algorithms properly necessitates the need for advanced instruction. But law librarians face an uphill battle when it comes to teaching advanced research concepts because "[t]oday's students arrive at law school often bereft of any research skills except the ability to 'Google.'"[187] "This is a challenge not unlike that faced by legal writing instructors who are expected to teach successful legal writing when they must first teach basic writing skills."[188]

¶80 The challenge of teaching students who lack even basic research skills is that "it [is] simply too convenient for people to follow the advice of an algorithm (or, too difficult to go beyond such advice)."[189] When students have relied on simple Google searches throughout their entire education without vetting the results to consider pitfalls like machine learning bias, it is of utmost importance that law librarians take the time to discuss these issues while teaching the foundations of legal research. "[A]lgorithms may lead to a loss in human judgment as people become reliant on the software to think for them."[190] And there is no greater threat than having machines create laws and govern populations without human understanding and oversight.

¶81 What we will see as search engines become more intelligent is that the systems will become even better at returning highly relevant results. "Seen through the prism of information theory, the legal information system will have improved its ability to communicate. Naturally, this improvement in capability should lead to changes in how the law is created and disseminated."[191] In turn, prospective lawyers will not have to spend a great deal of time looking through irrelevant results and instead will be able to focus on the implications of the results and parsing through how the results were generated.[192]

---

183. *Id.*

184. *Id.*

185. Valentine, *supra* note 130, at 204.

186. McGinnis & Wasick, *supra* note 6, at 1010.

187. Valentine, *supra* note 130, at 189.

188. *Id.*

189. Cecilia Mazanec, *Will Algorithms Erode Our Decision-Making Skills?*, NPR: All Tech Considered (Feb. 8, 2017), http://www.npr.org/sections/alltechconsidered/2017/02/08/514120713 /will-algorithms-erode-our-decision-making-skills.

190. *Id.*

191. McGinnis & Wasick, *supra* note 6, at 1022.

192. *Id.*

¶82 The responsibility to teach burgeoning technologies does not stop with law librarians. "Legal educators of the future will need to train new kinds of experts. Society will need a cadre of legal 'engineers' who can work with technologists to devise the new digital applications—hopefully while remembering that law has a public purpose."[193] Law schools should play an important role in educating nonlawyer developers to aid in best practices of legal algorithm creation. "Legal scholars are well positioned, if they can avoid the temptation to be rear-guard defenders of the old ways, to evaluate the risks and benefits of the new solutions, and to guide the debate on how they can be incorporated" into the practice of law.[194]

> It will take us some time to develop the wisdom and the ethics to understand and direct th[e] power [of algorithms]. In the meantime, we honestly don't know how well or safely [this power] is being applied. The first and most important step is to develop better . . . awareness of who, how, and where it is being applied.[195]

¶83 Moving forward, we must teach algorithmic literacy, transparency, and oversight concerns by which we provide education about how algorithms function in law.[196]

Conclusion

¶84 "It is premature to state categorically that computers will be used as aids in the process of legal reasoning, or even that they should be."[197] Those words, dating from the 1970s, no longer ring true. In 2018, we are closer to a time when computers will be used as aids in the process of legal reasoning, and it is beyond time to start considering how they should be.

¶85 In this interim period of NLP capability, when algorithms are used increasingly in the everyday practice of law, we must understand both the current limitations and the associated pitfalls. The strong PR campaigns of the latest and greatest technologies may exaggerate how well the technology performs given current NLP capabilities. While we can expect that PR folks will say certain things to sell a product, we cannot rely blindly on these claims or the products they describe.

¶86 Current and prospective lawyers must understand current computing capability to make an independent judgment regarding a system's abilities. They must consider algorithmic transparency and the associated machine learning bias that may be embedded into the results. Lawyers must also consider the ethical pitfalls such as the Duty of Technology Competence, supervisory requirements, and malpractice considerations.

¶87 And lawyers must do this while understanding that progress will come in fits and starts. After all, "[s]uch a system could be developed only to die of neglect; it could survive only in the cloisters of academia; it could become an occasional tool of some small or large number of lawyers; it could, conceivably, become a major influence in the practice of law."[198] This is where law librarians can be highly

---

193. Campbell, *supra* note 2, at 12.
194. *Id.* at 13.
195. Anderson & Rainie, *supra* note 155, at 16.
196. *Id.*
197. Buchanan & Headrick, *supra* note 127, at 60.
198. *Id.* at 61.

influential. Because law librarians are on the front lines of teaching legal research tools that increasingly rely on algorithms to perform the work, they are in the best position to teach prospective lawyers about the various issues surrounding the use of algorithms in law.

¶88 There are many conceivable futures for computers in law.

> As a profession, it is important that we don't identify with the pre-Gutenberg scribes. We are already in the business of using our expertise to help our constituents manage information overload. Watson and other augmented intelligence platforms are potentially powerful partners that can elevate and enhance our ability to manage the information tsunami pounding our desktops every day. It is clear that our role as experts in assessing and curating information quality will be more important than ever.[199]

¶89 As we law librarians consider the fate of law libraries in the Information Age and beyond, it is imperative that we continue to assess and instruct on information quality. "The fundamental aim of every law library ought to be to remind its patrons and constituents to dare to think otherwise—to see the law in its true, transformative essence. If law librarians do not play this important role, the battle may be lost entirely."[200]

---

199. O'Grady, *supra* note 86, at 21.

200. Paul D. Callister, *Law and Heidegger's Question Concerning Technology: A Prolegomenon to Future Law Librarianship*, 99 Law Libr. J. 285, 303, 2007 Law Libr. J. 285, ¶ 40.