Your Partner in the Profession | January 2020 • Vol. 89 • No. 1

THE JOURNAL OF THE KANSAS BAR ASSOCIATION

Illustration by Christian Book, KBA staff; Colored by Ryan Purcell, KBA staff

May It Please the Algorithm[1]

by Bob Lambrechts

Technology is rapidly advancing us to a crucial juncture in humanity's relationship with the law. In future disputes, machines may make life-and-death decisions all on their own. In 1970, the film "Colossus: The Forbin Project" brought to the theaters Dr. Charles A. Forbin's creation—a super computer designed to oversee and control America's huge military defense system. Not only will it control the nation's nuclear tipped missiles but it has limitless potential due to the sentience and artificial intelligence that Forbin embedded within the system. After launching a nuclear tipped missile at a Soviet oil field to convey a lethal lesson, Colossus tells Forbin that the world, under its absolute control, is now freed from war.

Litigation, an only marginally less so form of Armageddon, has rampant inefficiencies that make it nearly impossible to obtain an expeditious, on the merits resolution of even the most straightforward lawsuits. One commentator has noted:

> The recourse to legal actors and proceedings is costly, emotionally debilitating, and potentially counterproductive. The adversary system can be a hugely inefficient means of uncovering facts; its relentless formalities and ceaseless opportunities for splitting hairs are time consuming and expensive.[2]

Litigation is inefficient because it is ponderous and labored, which means that a client's financial and personnel resources can be redirected for extended periods of time. A large pending case can require a client to set aside substantial financial resources to address the litigation, to preserve numerous documents and utilize electronic search tools to scan a staggering number of company documents. Moreover, the client will have to assign personnel to address the demands of the litigation instead of having them focus on their revenue-generating roles within the company.

Whatever your attitude toward artificial intelligence, lawyers should count on the idea that technology will continue to change our profession and particularly how we practice litigation. It is difficult to think of a single area of modern technology that has not penetrated the law firm. Even if some lawyers fought against facsimile machines and computers at first, technology has made society more efficient and ironically may have increased the volume of legal services as lawyers began drafting their own documents and sending their own correspondence by electronic mail. The idea that a solo practitioner could function without an accountant and legal assistant would have been unthinkable fifty years ago, but can be standard procedure today.

This article will discuss the expanding role of artificial intelligence in the legal profession and the current and future roles of artificial intelligence in our legal system. I draw comparisons with the U.S. military's analysis of the role of artificial intelligence in war fighting. The basis for that comparison is that the Department of Defense is at the cutting edge of the development of this technology and highly ethical human judgment is critical to a morally acceptable outcome—just as with the legal profession. In addition, this article will discuss the prospect for using artificial intelligence in what might be considered the most hallowed roles within our legal system, that of the judge and jury, in their roles as the arbiters of justice.

Artificial Intelligence – Background and Projections

John McCarthy, a professor of computer science at Stanford, first conceived the term "artificial intelligence" in 1955.[3] In 1956, McCarthy invited a group of researchers from a multitude of disciplines, including language simulation, neural networks and complexity theory, to a summer workshop, the Dartmouth Summer Research Project on Artificial Intelligence, to discuss what would ultimately become the field of artificial intelligence.[4] It was evident many decades ago that electronic capacity and functionality were doubling approximately every eighteen months,[5] and the rate of improvement showed no signs of slowing down. In fact, experts predict that spending on artificial intelligence by companies will grow from $37.5 billion in 2019 to nearly $98 billion in 2023, a compound annual growth rate of 28.4 percent during the period between 2018 and 2023.[6]

The Dartmouth conference was one of the first serious attempts to consider the consequences of this exponential curve. Many attendees came away from the conference convinced that continued advancements in electronic speed, capacity, and software programming would lead to the point where computers would someday have the resources to be as intelligent as human beings.[7]

Artificial intelligence is the science and engineering of making intelligent machines. It is not a single technology but is comprised of related and often-connected technologies that work together to supply "human-like" responses and reasoning. Also referred to as "cognitive technologies," artificial intelligence comprises, among other things, the technologies of deep learning, natural language processing, machine vision, speech recognition and expert systems.[8] Among these, deep learning is the most transformative and is the core of what is considered modern artificial intelligence. Deep learning utilizes neural networks, a computer system modeled after the human brain and nervous system, that learn from large amounts of data.[9] This is akin to how we learn from experience. The deep learning algorithm would perform a task repeatedly, each time tweaking it a little to improve the outcome.[10] We refer to "deep learning" because the neural networks have various (deep) layers that enable learning. "Just about any problem that requires 'thought' to figure out is a problem deep learning can learn to solve."[11]

It is estimated that every day we generate a mind-boggling 2.5 quintillion bytes with 90 percent of all data today created in the last two years.[12] Since deep-learning algorithms require enormous amounts of data to learn from, this increase in data creation is one reason that deep learning capabilities have grown in recent years. In addition to more data creation, deep learning algorithms benefit from the more robust computing power that is available today. It is computing capacity that makes deep learning possible. The typical human brain is made of an estimated 86 billion interconnected brain cells, or neurons.[13] To provide a sense of the advances that are being made in this area, in the summer of 2019 Intel made considerable progress toward a digital equivalent of the human brain by building a computer system with 8 million digital neurons and has the goal of reaching 100 million by late 2019.[14]

Artificial neural networks seek to simulate these biological networks and get computers to act like interconnected brain cells, so that they can learn and make decisions in a more humanlike manner.[15] Discrete areas of the human brain process information differently, and these parts of the brain are arranged in a hierarchical fashion.[16] As information enters the brain, "each level of neurons processes the information, provides insight and passes the information to the next, more senior layer."[17]

With deep learning, the computer trains itself to process and learn from data. According to Ray Kurtzweil, an American inventor, futurist and director of engineering at Google, by 2045, computers utilizing artificial intelligence will surpass human intelligence.[18] He describes uploading as a process of "scanning all of the salient details (of a human brain) and then reinstantiating those details into a suitably powerful computational substrate.[19] This process would capture a person's entire personality, memory, skills and history."[20]

Deep learning is a method for software to learn by trial and error at a pace limited only by computer processing power and cloud storage.[21] Using unstructured data[22] (80 percent of all the data that exists is unstructured)[23] and operating without the need for explicit, step-by-step instructions, deep learning systems iteratively generate solutions.[24] The outcome from many deep learning iterations is a digital neural network considered comparable to how humans think, which establishes patterns, relationships and connections within data that is otherwise unstructured data.

Now, "new machine learning approaches literally have the machines learn on their own things that we don't know how to explain."[25] The machines learn patterns, correlations and rules, sometimes the ones that humans use to accomplish the task but other times ones that humans cannot discern.[26] Indeed, many times the programmer cannot account for how the machine came to a particular result, even if the result is correct.[27] Tasks that were once impossible to automate are now on par with human experts, including not only facial recognition,[28] but also skin cancer detection[29] and some types of language translation.[30] IBM's Watson, for example, analyzed questions and content comprehensively and quickly and eventually won "Jeopardy!" against former champions.[31] Reinforcement learning, a category of machine learning, entails experimentation.[32] Reinforcement learning is already prevalent in some forms of artificial intelligence. A computer developed by a subsidiary of Alphabet learned and mastered Go, a notoriously complicated board game, and eventually

beat one of the world's best human players.[33] With reinforcement learning, "the neural network is reinforced for positive results, and punished for a negative result, forcing the neural network to learn over time."[34]

Autonomous Weapon Systems and the U.S. Military

While to some it may appear to be a non-analogous leap to commingle the discussion of artificial intelligence in a legal context with the discussion of the same technology in a military context, there are concerns that both systems face with the use of this technology, and in particular the military's experience, struggles and constraints.

In the 115th Congress, thirty-nine bills included the phrase "artificial intelligence" in the text of the bill[35] and incorporated an often-cited classification scheme that categorizes artificial intelligence systems as designed to think rationally, act rationally, think like humans, or act like humans.[36] Several of these bills were enacted into law.[37] These classifications were broadly incorporated into the first statutory definition of artificial intelligence, included in the John S. McCain National Defense Authorization Act for Fiscal Year 2019 (P.L. 115-232). Section 238 provides that the term "artificial intelligence" includes:

> (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
>
> (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication or physical action.
>
> (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
>
> (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task.
>
> (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot, that achieves goals using perception, planning, reasoning, learning, communicating, decision-making and acting.

Department of Defense (DOD) Directive 3000.09, Change 1, dated May 8, 2017, titled "Autonomy in Weapons Systems," provides that autonomous and semi-autonomous weapon systems are to be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force.[38] Additionally, the DOD does not currently have an autonomous weapon system that can search for, identify, track, select and engage targets independent of a human operator's input.[39]

The DOD directive also stipulates that "autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force," precluding the development of fully autonomous weapons systems.[40] This reluctance to pursue fully autonomous weapons systems was further emphasized during 2017 testimony to the Senate Armed Services Committee, when then-Vice Chairman of the Joint Chiefs of Staff General Paul Selva stated, "I am an advocate for keeping the restriction, because we take our values to war.... I do not think it is reasonable for us to put robots in charge of whether or not we take a human life."[41]

The standard for autonomous weapon systems' compliance with the laws of war should arguably not be whether they are able to make unflawed decisions, but whether they are able to follow the principles of proportionality, military necessity and distinction, at least as well as human operators.[42] "It must be emphasized that as a matter of law, more may not be asked of autonomous weapon systems than of human-operated systems."[43] With increasingly sophisticated sensors linked to advanced artificial intelligence, in the not too distant future, autonomous systems may be capable of distinguishing between civilians and combatants at a level comparable to a human operator in at least some battlefield environments. In that instant, it will be the responsibility of commanders to ensure that any autonomous weapon systems used is capable of distinguishing between civilians and combatants in the environment in which they are deployed.[44]

Opponents also argue that human compassion and other emotions are necessary to ethical war-fighting.[45] Human empathy, some argue, helps soldiers to assess the objectives of potential human targets to discern whether they really pose a threat.[46] Machines may possibly never be programmable to effectively emulate empathy.[47] "On the other hand, proponents of such systems argue that human emotions—fear, anger and the instinct for self-preservation—may lead to adverse consequences on the battlefield. Robots, they posit, may not be subject to human errors or unlawful behavior induced by human emotions."[48]

War fighters must be mindful of over-dependency upon situational data provided by an autonomous system in order to avoid excessive reliance by the combatants upon the "judgment" of the autonomous system rather than their own seasoned judgment. This type of over-reliance could lead to the phenomenon of "automation bias."[49] Taken to a logical conclusion, this reliance upon the assessment provided by the autonomous system can lead to a psychological detachment from the consequences of the delivery of weapons systems and make killing too remote for soldiers.[50]

Any standardized position on the use of artificial intelligence for battlefield decisions should ultimately be informed by verifiable scientific data on the benefits and drawbacks of human as opposed to machine decision-makers.[51] If it turns out that the machine is less prone to mistakes and less likely than humans to be influenced by, for example, fear or hatred, then the war fighter should be required to use artificial intelligence in decision-making in order to reduce harm to civilians.[52] "Failure to use such technology may be regarded as failure to apply a reasonable precaution."[53]

When there are large gaps between the data gathering and analysis capacity of machines and humans, human control provides less value. Importantly, it is these gaps—which are bound to grow with the utilization of advanced and abundant battlefield sensing systems—that will lead humans away from being decision-makers in many war fighting situations and effectively being repositioned as a check on artificial intelligence decision-making power.[54]

Artificial Justice

As with the military utilization detailed above, artificial intelligence is permeating numerous aspects of the practice of law. Much like the military, the U.S. legal system is wary of the prospect of turning over total control of adjudication to machines. Nonetheless, there is a steady erosion of tasks that at one time were considered too challenging for computational involvement. For example, many legal research services such as CARA, Clerk, EVA, and vLex now include brief-evaluation tools that use artificial intelligence to analyze a brief, whether for a client or from an opposing party.[55]

These services look at factors such as the procedural posture of the case, the pattern of citations, and even which citations may be missing.[56] They can evaluate strengths or weaknesses of a brief or pleading based on which claims are made or omitted. Researchers at LegalMation have created document-automation tools that ingest complaints and with artificial intelligence create the first draft of responsive pleadings, even though for a limited number of causes of action and in a small number of jurisdictions.[57]

Artificial intelligence capabilities are growing and will undoubtedly be deployed into activities that are presently seen as the sole purview of the human jurist. The role of artificial intelligence systems in litigation includes a range of possibilities from the increasing use of technology in legal and judicial processes prior to trial to having some involvement in court proceedings. Even before a case is docketed, artificial intelligence may already affect how cases are prepared and presented to the court.[58]

The very role of a trial in our modern American justice system is to provide a fair process in which to definitively resolve an otherwise intractable dispute. A basic premise of our legal system is that if the trial is procedurally fair, the outcome of the process is presumed to be correct. If people consider that they have been treated fairly, they are more likely to accept a decision and outcome.[59] The presumption is intentionally designed to prevent the potential for endless re-litigation of cases. The American legal system does, of course, recognize the prospect of trial error, and includes a sophisticated appeals process to correct errors.

Importantly in this analysis, the Supreme Court has opined that jurors "are presumed to be fitted" for "[d]etermining the weight and credibility of witness testimony" by "their natural intelligence and their practical knowledge of men and the ways of men,"[60] and that a "fundamental premise of our criminal trial system is that 'the jury is the lie detector.'"[61] But does the public necessarily believe in the jury's ability to determine fabrications of the truth? The myth that "lie detecting is what our juries do best"[62] may no longer be credible in the age of DNA-based convictions and exonerations. Nor are jurors particularly skilled at determining when witnesses are credible but mistaken, as empirical studies on eyewitness identification testimony have shown.[63]

According to a study performed in 2000-2001 by Northwestern University statistician Dr. Bruce Spencer, juries wrongfully convict in about thirteen percent of cases.[64] When they err, they do so more asymmetrically; they are more likely to convict an innocent person than to acquit a guilty person.[65] It is only to be expected that any system relying exclusively on human judgment will make mistakes. The number of convicted felons who have been exonerated by improved DNA and other forensic tests is a disturbing reminder on the imperfections of our justice system and ourselves.

Judges and juries bring different skill sets and life experiences to their roles.[66] Judges are not only legal experts, but also have the advantage of becoming very familiar with the evidence that is presented in similar cases over an extended period. Juries on the other hand, have the benefit of pooling their education and experience, which maximizes their collective recollection and comprehension of the evidence and minimizes the possibility that biases held by any one juror will inappropriately skew the jury's interpretation of evidence.

Enter artificial intelligence and the impacts of that technology. It may begin even with influencing which cases get before a judge, as predictive coding developments enable predictions to be made as to the outcome of litigation.[67] In Mexico, the Expertius system is advising judges and clerks "upon the determination of whether the plaintiff is or is not eligible for granting him/her a pension."[68] In the United States, predictive coding has been used to help determine whether recidivism is more likely in criminal matters and to assist in making decisions about sentencing.[69]

Judicial Temperament

"Judicial responsiveness requires judges to act from the perspective of conscious legal rationality and also with intuition, empathy and compassion."[70] The function of the human judge is not to crunch data. Some would undoubtedly argue that the attribute of the humaneness of the judge must be infused within the technology so that it plays a principled role in advancing a responsive justice system. At least one MIT professor of artificial intelligence has postulated that due to the interplay of the brain's frontal lobe and limbic system (the subcortical areas that play a critical role in pattern recognition of sound, vision, and smell), our ability to reason and to weigh the value of information depends in part on our ability to feel emotion.[71]

Many artificial intelligence experts believe the converse is true and see no merit in the role of emotion, preferring to build systems that rely solely on rules.[72] Stanford computer science professor John McCarthy argues that emotion should not be a consideration in computing, that emotion is not essential to intelligence, and can be problematic.[73] "The goal is to build machines that apply certain human values and principles in decision-making."[74] Computational cognitive modeling,[75] for example in the area of contractual approach to ethics[76] is being used to describe principles used in decision making and to determine how human minds apply those rules.

Affective Processing

Assuming arguendo that emotional artificial intelligence is central to the transition to machine adjudication, is there a path forward or is it simply beyond the pale of human endeavor to infuse such capabilities into a machine? As noted above, we seek for our judges to act with intuition, empathy and compassion. Artificial intelligence and neuroscience researchers agree that artificial intelligence, at present, does not have its own emotions.[77] Artificial intelligence, however, is rapidly evolving and advances in "affective processing,"—a field of study and development of systems and devices that can recognize, interpret, process and simulate human affects—may yield a form of machine empathy.[78] One critical challenge for researchers is to develop the ability to simulate empathy. Ideally, in a judicial setting, the machine should interpret the emotional state of the participants and adapt its behavior to them, giving an appropriate response to those emotions.

Building a machine that can perceive emotional signals is distinct from teaching a machine to interpret them. Expressing emotion is yet another separate task.[79] "In a machine . . . you can decouple capabilities – train it to recognize anger but give it no feelings. And you can go pretty far with this, making it perceive or even express emotions but without the actual feelings."[80] The task of designing an integrated artificial moral decision-making system is complicated by our still-evolving understanding of how human moral decision-making actually works.

Selecting the morally critical features of a situation is a complex undertaking that may proceed largely unconsciously, and is likely supported by a combination of theoretical reasoning and emotional intelligence.[81] Emotions also arguably provide important channels for acquiring information relevant to moral decision-making.[82] Some would assert that the human ability to assess the emotional state of others is essential to our ability to respond in an appropriate manner.

Bias – Human and Machine

The growing use of artificial intelligence in sensitive areas, including for hiring, criminal justice, and healthcare, has stirred a debate about bias and fairness. "Some researchers have highlighted how judges' decisions can be unconsciously influenced by their own personal characteristics."[83]

At least one author has argued that the motivating impulse which leads a judge to his decision is his "intuitive sense of what is right or wrong in the particular case."[84] Once the decision is rendered, the judge will work aggressively to justify his or her decision within his or her own mind and to withstand criticism from peers.[85] The judicial hunch "is a composite reaction to a multitude of responses to the stimuli set up by witnesses—stimuli which encounter the judge's biases, stereotypes, preconceptions and the like."[86] It appears that bias, whether human or algorithmic in nature is, and will continue to be, a source of concern.

In the world of artificial intelligence, it is generally the data as opposed to the algorithm itself that is most often the main source of the issue.[87] Bad data can contain implicit racial, gender or ideological biases.[88] "Models may be trained on data containing human decisions or on data that reflect second-order effects[89] of societal or historical inequities." As computer scientists work to develop artificial intelligence systems that can be trusted, it is critical to develop and train these systems with data that is unbiased, and to develop algorithms that can be easily explained. More than 180 human biases have been defined and classified, any one of which can affect how we make decisions.[90]

There are two principal mechanisms by which bias shows up in training data. The first is that the data reflects existing prejudices and the second, it is unrepresentative of reality.[91] The latter situation might occur, for example, if a deep-learning algorithm is fed more photos of dogs than cats. The resulting animal recognition system would understandably be less adept at recognizing cats.

The first path for introducing bias can be illustrated with Amazon's efforts in 2014 to build a computer program to review job applicants' resumes with the aim of mechanizing the search for exceptional talent.[92] By 2015, the company re-

cognized that the algorithm in concert with the data was not rating candidates for software developer jobs and other technical posts in a gender-neutral manner. Amazon's computer models were trained to screen applicants by observing historical patterns in resumes submitted to the company.[93] Most of the resumes were provided by men which is a reflection of the male dominance within the tech industry.[94] The Amazon system taught itself that male candidates were more desirable. It downgraded resumes that, for example, included the word "women," as in "women's chess club captain."[95]

It is also possible to introduce bias during the data preparation stage, which involves selecting which attributes the algorithm is to consider. In the case of modeling soil quality, for example, an "attribute" could be the concentration of certain nutrients, hydraulic conductivity or moisture content. In the case of Amazon's recruiting tool, an "attribute" could be the candidate's gender, education level or years of experience. This is what is often referred to as the "art" of deep learning: choosing which attributes to consider or ignore can significantly influence the model's prediction accuracy. But while its impact on accuracy is easy to measure, its impact on the model's bias is not.

Bias can also arise based on the biases of the users driving the interaction. A clear example of this bias is Microsoft's Tay (acronym for "thinking about you"), an artificial intelligence chatbot[96] that was originally released by Microsoft via Twitter on March 23, 2016.[97] Tay caused considerable controversy when the bot began to tweet all sorts of misogynistic and racist remarks, causing Microsoft to shut it down in less than 24 hours after its launch.[98] According to Microsoft, this was caused when people tweeted the service and the bot made replies based on its interactions with the people on Twitter.[99]

As intelligent systems are built that make decisions with and learn from human partners, the same sort of defect might arise in more problematic circumstances, including within the legal system. One solution may be to partner individuals with the intelligent system that can guide them over time. What was learned from the Tay experience was that such systems will infuse the biases of people interacting with the system, thereby reflecting the opinions of the people who train them.

Identifying, mitigating, and hopefully someday eliminating bias in artificial intelligence systems is essential to building trust between humans and machines that learn. IBM Researcher Francesca Rossi opines that "As AI systems find, understand and point out human inconsistencies in decision making, they could also reveal ways in which we demonstrate partial, parochial, and cognitive biases, leading us to adopt more impartial or egalitarian views."[100] "In the process of recognizing our bias and teaching machines about our common values, we may improve more than AI. We might just improve ourselves."[101]

In summary, there are many complicated hurdles to overcome with the development of artificial intelligence systems before they attain the capability of assuming the role of fair jurists. Nonetheless, it is this author's opinion that with the passage of time, we shall see these systems advance into roles that today seem quite unimaginable. ■

,

About the Author

Bob Lambrechts is a partner in the Overland Park office of Lathrop GPM, LLP, where he practices in the areas of patent and environmental law. He formerly served with the U.S. Environmental Protection Agency in the Clean Air Division and as a special assistant to the Regional Administrator. In addition to a law degree from St. Louis University, Mr. Lambrechts earned a Bachelor's and Master's Degree in Mechanical and Aerospace Engineering from the University of Missouri at Columbia and prior to practicing law, he served as a robotics engineer with the Bendix Corporation in Kansas City. As avocations, Mr. Lambrechts served as an Engineering Duty Officer in the U.S. Navy Reserves for 28 years and since 1993, as an adjunct faculty member for the University of Missouri at Kansas City Schools of Law and Engineering.

blambrechts@lathropgage.com

1. "It is often said that *May it please the Court* is an obligatory phrase at the outset of an oral argument—and that any other opener suggests the oral advocate is unknowledgeable or inexperienced." Bryan A. Garner, *What judges really think about the phrase 'May it please the Court?'*, ABA Journal, Bryan Garner on Words, April 1, 2013, http://www.abajournal.com/magazine/article/what_judges_really_think_about_the_phrase_may_it_please_the_court.
2. Craig C. Martin, *Avoiding the Inefficiency of Litigation*, Pretrial Practice & Discovery, Vol. 15, Number 3, Spring 2007, ABA, https://jenner.com/system/assets/publications/2105/original/PP_D_Martin_Spring07.pdf?1315481513.
3. John McCarthy, Computer History Museum, https://www.computerhistory.org/fellowawards/hall/john-mccarthy/.
4. Bernard Marr, Forbes, *The Key Definitions of Artificial Intelligence (AI) that Explain Its Importance*, Feb14, 2018 https://www.forbes.com/sites/bernardmarr/2018/02/14/the-key-definitions-of-artificial-intelligence-ai-that-explain-its-importance/#72321b8c4f5d.
5. In 1965, Gordon Moore, a scientist at Intel, made a prediction based on his observation that the number of transistors per square inch on integrated circuits had doubled every year since their invention. He predicted that the trend will continue, and growth in computer power will double roughly every two years while the cost of that computing power will go down. When coupled with the ever-lower cost of storing electronic data, you have the basis for the rapid rise in AI capabilities and availability.
6. Richard Seeley, Application Development Trends Magazine, Global Spending on AI Systems to Hit $98 Billion by 2023, September 4, 2019, https://adtmag.com/articles/2019/09/04/ai-spending.aspx.
7. Bernard Marr, Forbes, *The Key Definitions of Artificial Intelligence (AI) that Explain Its Importance*, Feb14, 2018 https://www.forbes.com/sites/bernardmarr/2018/02/14/the-key-definitions-of-artificial-intelligence-ai-that-explain-its-importance/#72321b8c4f5d.

8. Michael Mills, *Artificial Intelligence in Law: The State of Play 2016*, https://www.neotalogic.com/wp-content/uploads/2016/04/Artificial-Intelligence-in-Law-The-State-of-Play-2016.pdf.

9. Jonas DeMuro, *What is a neural network?* World of Tech, August 11, 2018, https://www.techradar.com/news/what-is-a-neural-network.

10. *Id.*

11. Bernard Marr, *What is Deep Learning AI? A simple guide with 8 practical examples,* 2019, https://bernardmarr.com/default.asp?contentID=1572

12. Data Never Sleeps 5.0, https://web-assets.domo.com/blog/wp-content/uploads/2017/07/17_domo_data-never-sleeps-5-01.png.

13. Stephen Shankland, Intel packs 8 million digital neurons onto its brain-like computer, https://www.cnet.com/news/intel-packs-8-million-digital-neurons-onto-brain-like-pohoiki-beach-computer-loihi-chips/.

14. *Id.*

15. Bernard Marr, *Deep Learning Vs Neural Networks – What's the Difference?* https://bernardmarr.com/default.asp?contentID=1789.

16. *Id.*

17. *Id.*

18. *Kate Kershner, The Singularity by 2045, Plus 6 other Ray Kurzweil Predictions, Howstuffworks, March 22, 2017*, https://electronics.howstuffworks.com/future-tech/singularity-2045-ray-kurzweil-predictions.htm.

19. Ray Kurzweil, *The Singularity is Near: When Humans Transcend Biology* (New York, 2005), p. 374.

20. *Id.*

21. Robert D. Hof, *Deep Learning*, MIT Tech. Rev. (2013).

22. Unstructured data is information that either does not have a pre-defined data model or is not organized in a pre-defined manner. Unstructured information is typically text-heavy, but may contain data such as dates, numbers, and facts as well.

23. Devin Pickell, *Structured vs Unstructured Data – What's the Difference*, November 16, 2018, https://learn.g2.com/structured-vs-unstructured-data.

24. "Feedback neural networks contain cycles. Signals travel in both directions by introducing loops in the network. The feedback cycles can cause the network's behavior change over time based on its input. Feedback neural network also known as recurrent neural networks." Avinash Navlani, *Neural Network Models in R,* Datacamp, January 18, 2019, https://www.datacamp.com/community/tutorials/neural-network-models-r.

25. *How AI Is Already Changing Business,* HBR IDEACAST (July 20, 2017), https://hbr.org/ideacast/2017/07/how-ai-is-already-changing-business.

26. Heather Murphy, *Why Stanford Researchers Tried to Create a 'Gaydar' Machine,* N.Y.TIMES (Oct. 9, 2017), https://www.nytimes.com/2017/10/09/science/stanford-sexual-orientation-study.html .

27. Andreas Holzinger *et al.*, *What Do We Need to Build Explainable AI Systems for the Medical Domain?,* ARXIV.ORG (2017), https://arxiv.org/pdf/1712.09923.pdf ("However, the central problem of such models is that they are regarded as black-box models and even if we understand the underlying mathematical principles of such models they lack an explicit declarative knowledge representation, hence we have difficulty in generating the underlying explanatory structures.")

28. *See generally* Will Knight, *Paying with Your Face,* MIT TECH. REV. (Mar.-Apr. 2017), https://www.technologyreview.com/s/603494/10-breakthrough-technologies-2017-paying-with-your-face/ (detailing how researchers have shown their programs rival most humans in ability to recognize faces).

29. Andre Esteva et al., *Dermatologist-level Classification of Skin Cancer with Deep Neural Networks,* 542 NATURE 115, 118 (2017), https://www.nature.com/articles/nature21056

30. Jason Brownlee, *How to Develop a Neural Machine Translation System from Scratch,* Machine Learning Mastery, August 7, 2019, https://machinelearningmastery.com/develop-neural-machine-translation-system-keras/

31. The DeepQA Research Team, IBM RES., http://researcher.watson.ibm.com/researcher/viewgroup.php?id=2099 ("[Watson] analyzes natural language questions and content well enough and fast enough to compete and win against champion players at *Jeopardy!*")

32. Will Knight, *Reinforcement Learning,* MASS. INST. TECH., TECH. REV. (2017), https://www.technologyreview.com/s/603501/10-breakthrough-technologies-2017-reinforcement-learning [https://perma.cc/M77M-RBCG].

33. Matthew Hutson, *This computer program can beat humans at Go – with no human instruction,* Science, October 18, 2017, https://www.sciencemag.org/news/2017/10/computer-program-can-beat-humans-go-no-human-instruction.

34. James Furbush, Cognex Blogs, *How a neural network helps manufacturing inspection, September 17, 2019,* https://www.cognex.com/blogs/deep-learning/what-is-a-neural-network.

35. Results of bill text only search on Congress.gov using the phrase "artificial intelligence" on November 13, 2019.

36. Peter Norvig and Stuart J. Russell, *Introduction to AI: A Modern Approach*, Chapter 1, https://people.eecs.berkeley.edu/~russell/intro.html.

37. In addition to the John S. McCain National Defense Authorization Act for Fiscal Year 2019, see FAA Reauthorization Act of 2018, Pub. L. 115-254, § 548, 132 Stat. 3186, https://www.congress.gov/115/bills/hr302/BILLS-115hr302enr.pdf; Fixing America's Surface Transportation Act, Pub. L. 114-94, § 6004, 129 Stat. 1312, 1562 (2015), https://www.congress.gov/114/plaws/publ94/PLAW-114publ94.pdf.

38. Kelley M. Sayler, Defense Primer: U.S. Policy on Lethal Autonomous Weapon Systems, March 27, 2019, Congressional Research Service; https://crsreports.congress.gov/product/pdf/IF/IF11150.

39. Zachary Fryer-Biggs, The Pentagon Plans To Spend $2 Billion To Help Inject More Artificial Intelligence Into Its Weaponry, October 29, 2018, https://publicintegrity.org/national-security/the-pentagon-plans-to-spend-2-billion-to-help-inject-more-artificial-intelligence-into-its-weaponry/.

40. *Id.*, p. 2.

41. Senate Committee on Armed Services, Hearing to Consider the Nomination of General Paul J. Selva, USAF, for Reappointment to the Grade of General and Reappointment to be Vice Chairman of the Joint Chiefs of Staff, 115th Congress, 1st Session, July 18, 2017.

42. Schmitt, Michael N., and Jeffrey S. Thurnher. *Out of the Loop: Autonomous Weapon Systems and the Law of Armed Conflict.* Harvard National Security Law Journal, vol. 4., p. 243 (2013). https://harvardnsj.org.

43. *Id.* at 247.

44. *Id.* p. 252-53, 278.

45. Bonnie Docherty, Human Rights Watch, *et al.*, *Losing Humanity: The Case Against Killer Robots*, November 19, 2012, https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots.

46. Congressional Research Service, *U.S. Ground Forces Robotics and Autonomous Systems (RAS) and Artificial Intelligence (AI): Considerations for Congress,* p. 34, November 20, 2018, https://fas.org/sgp/crs/weapons/R45392.pdf.

47. *Id.*

48. *Id.*

49. Heko Scheltema, *Lethal Automated Robotic Systems and Automation Bias*, Blog of the European Journal of International Law, June 11, 2015, https://www.ejiltalk.org/lethal-automated-robotic-systems-and-automation-bias/.

50. *Id.*

51. Yuval Shany, *Expert views on the frontiers of artificial intelligence and conflict*, March 19, 2019, https://blogs.icrc.org/law-and-policy/2019/03/19/expert-views-frontiers-artificial-intelligence-conflict/.

52. *Id.*

53. *Id.*

54. *Id.*

55. Bob Ambrogi, *Vincent Joins CARA, EVA and Clerk as the Latest AI-Driven Research Assistant,* Lawsites (Sept. 20, 2018), https://www.

lawsitesblog.com/2018/09/vincent-joins-cara-eva-clerk-latest-ai-driven-research-assistant.html.

56. *Id.*

57. Legalmation Platform Overview, https://www.legalmation.com.

58. Lopez, I. (2016) *The early years begin for AI's transformation of law,* Legaltech News, October 5, 2016. http://www.legaltechnews.com/id=1202769286334/The-Early-Years-Begin-for-AIs-Transformation-of-Law?cmp=share_twitter&slreturn=20160912054113.

59. Tom R. Tyler, *Social Justice: Outcome and Procedure*, International Journal of Psychology, New York University, 2000, volume 35 (2), p. 117-125, https://onlinelibrary.wiley.com/doi/pdf/10.1080/002075900399411.

60. *United States v. Scheffer*, 523 U.S. 303, 313 (1998) (quoting *Aetna Life Ins. Co. v. Ward*, 140 U.S. 76, 88 (1891)).

61. *Id.* at 313 (quoting *United States v. Barnard*, 490 F.2d 907, 912 (9th Cir. 1973).

62. *See* George Fisher, *The Jury's Rise as Lie Detector*, 107 Yale L.J. 575, 578–79 (1997).

63. Steven I. Frieland, *On Common Sense and the Evaluation of Witness Credibility,* Case Western Reserve University School of Law, Vol. 40, 1989, Issue 1, p. 167, https://scholarlycommons.law.case.edu/cgi/viewcontent.cgi?article=1865&context=caselrev.

64. Bruce Spencer, *Estimating the Accuracy of Jury Verdicts*, Northwestern University Institute for Policy Research, June 18, 2007, https://www.ipr.northwestern.edu/our-work/working-papers/2006/ipr-wp-06-05.html.

65. *Id.*

66. Toby S. Goldbach and Valerie P. Hans, *Juries, Lay Judges, and Trials* (2014). Cornell Law Faculty Working Papers, Paper 122. http://scholarship.law.cornell.edu/clsops_paper/122.

67. Schubarth C. (2016) *Y combinator startup uses big data to invest in civil lawsuits.* Silicon Valley Bus. J., 24 August. http://www.bizjournals.com/sanjose/blog/techflash/2016/08/y-combinator-startup-uses-big-data-to-invest-in.html.

68. Carniero D, Novais P., Andrade F, Zeleznikow J., Neves J. (2014) *Online Dispute Resolution: An Artificial Intelligence Perspective.* Artif Intell Rev 41(2):211-240, https://ddd.uab.cat/pub/artpub/2014/202519/air_a2014v41n2p211iENG.pdf.

69. Liptak A. (2017) *Sent to prison by a software program's secret algorithms.* New York Times (online), May 1. https://www.nytimes.com/2017/05/01/us/politics/sent-to-prison-by-a-software-programs-secret-algorithms.html?smid=tw-share&_r=0.

70. Tania Sourdin and Richard Cornes, *Do Judges Need to be Human? The Implications of Technology for Responsive Judging*, 2018.

71. David Diamond, *The Love Machine*, Wired Magazine, December 1, 2003, https://www.wired.com/2003/12/love/.

72. *Id.*

73. *Id.*

74. Aitrends, IBM Research: *Many AI Systems are Trained Using Biased Data*, February 1, 2019, https://www.aitrends.com/ethics-and-social-issues/ibm-research-many-ai-systems-are-trained-using-biased-data/.

75. Computational cognitive modeling aims to understand behavioral data and the mind and brain, more generally, by building computational models of the cognitive processes that produce the data.

76. "Social contract theory says that people live together in society in accordance with an agreement that establishes moral and political rules of behavior. Some people believe that if we live according to a social contract, we can live morally by our own choice and not because a divine being requires it. . . . Social contracts can be explicit, such as laws, or implicit, such as raising one's hand in class to speak. The U.S. Constitution is often cited as an explicit example of part of America's social contract. It sets out what the government can and cannot do. People who choose to live in America agree to be governed by the moral and political obligations outlined in the Constitution's social contract." Social Contract Theory, *Ethics Unwrapped*, McCombs School of Business, https://ethicsunwrapped.utexas.edu/glossary/social-contract-theory.

77. Herve-Victor Fomen, *Can Artificial Intelligence become emotional?* November 2018, https://www.leyton.com/blog/?p=3193-can-artificial-intelligence-become-emotional.

78. Affective computing originated with Rosalind Picard's 1995 paper, *Affective Computing*, MIT Technical Report #321, 1995. A motivation for the research is the ability to simulate empathy. The machine should interpret the emotional state of humans and adapt its behavior to them, giving an appropriate response to those emotions. https://affect.media.mit.edu/pdfs/95.picard.pdf.

79. David Diamond, *The Love Machine,* Wired Magazine, December 1, 2003, https://www.wired.com/2003/12/love/.

80. *Id.*

81. *Moral Reasoning*, Stanford Encyclopedia of Philosophy, August 27, 2018, https://plato.stanford.edu/entries/reasoning-moral/.

82. *Id.*

83. *Tackling bias in artificial intelligence (and in humans),* McKinsey & Company, June 2019, https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans.

84. Timothy J. Capurso, (1998) *How Judges Judge: Theories on Judicial Decision Making*, University of Baltimore Law Forum: Vol. 29: No. 1, Article 2, available at http://scholarworks.lawubalt.edu/lf/vol29/iss1/2.

85. *Id.*

86. *Id.*

87. Jake Silberg and James Manyika, *Tackling bias in artificial intelligence (and in humans)*, McKinsey Global Institute, June 2019; https://www.mckinsey.com/featured-insights/artificial-intelligence/tackling-bias-in-artificial-intelligence-and-in-humans.

88. *Id.*

89. *Id.* Second-Order Effect refers to the idea that every action has a consequence, and each consequence has a subsequent consequence. In other words, this means that a single decision can initiate a series of cause-and-effects, something which we might not have knowledge or control of.

90. Wikipedia, *List of cognitive biases*, https://en.wikipedia.org/wiki/List_of_cognitive_biases.

91. Rashmi Margani, *What is AI Bias?, How it affects?* Medium.com, April 24, 2019, https://medium.com/@rashmi.margani/what-is-ai-bias-how-it-affects-65ff2cdcee70.

92. Jeffrey Dastin, *Amazon scraps secret AI recruiting tool that showed bias against women,* Reuters, October 9, 2018, https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G.

93. *Id.*

94. James Vincent, *Amazon reportedly scraps internal AI recruiting tool that was biased against women,* The Verge, October 10, 2018, https://www.theverge.com/2018/10/10/17958784/ai-recruiting-tool-bias-amazon-report.

95. *Id.*

96. A chatbot is a computer program that simulates human conversation through voice commands or text chats or both. Chatbot, short for chatterbot, is an Artificial Intelligence (AI) feature that can be embedded and used through any major messaging application.

97. James Vincent, *Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day*, The Verge, March 24, 2016, https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist.

98. *Id.*

99. *Id.*

100. Hannah Herman, *AI is Biased. Here's How We Can Fix That,* Workato, November 27, 2018, https://blog.workato.com/2018/11/artificial-intelligence-bias/#.Xc8fXFdKgdU.

101. *Id.*