SQUIRE PATTON BOGGS

Artificial Intelligence: A Basic Primer

In the summer of 1956, a Dartmouth College mathematics professor organized the Dartmouth Summer Research Project on Artificial Intelligence in Hanover, New Hampshire. In his proposal to organize the conference, Professor John McCarthy indicated that its purpose was "to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can, in principle, be so precisely described that a machine can be made to simulate it."

Over half a century later, we regularly read and hear about Artificial Intelligence (AI), but we are challenged to predict with confidence its impact on society. Google's CEO, Sundar Pichai, recently predicted that it would be "more profound than…electricity or fire." We know from personal experience that AI can enrich our lives, even if we don't know how the technology does it. As we benefit from advances in AI, we are confronted with doomsday predictions about robots someday ruling the earth if we don't control AI's ascendancy. Rather than getting carried away, we think it most helpful to define AI, to describe where it is today, and then to offer our thoughts about how it will change everything.

Although the term "Artificial Intelligence" covers a wide range of computerized activity, we think of it as the ability of a machine, in some way, to solve problems by performing cognitive functions associated with the human mind, such as perceiving, learning, planning, reasoning, and interacting with our physical environment. In performing or mimicking these cognitive functions, AI can do smart things in real time on a scale that humans cannot, such as finding patterns in trillions of data points and predicting resulting outcomes or actions.

When we use this definition, we have in mind "narrow" or "weak" AI, that is, the ability of a machine to focus on one narrow task, performing it at unfathomable speed with exceptional accuracy. By contrast, "artificial general intelligence" (AGI) or "strong" AI is the ability of a machine to truly think like a human, to reason, to understand cause and effect – potentially with unpredictable, uncontrolled consequences to society.

Artificial Intelligence is extremely powerful applied technology, with the potential to transform many aspects of human existence (including the development and delivery of legal advice). But it is not "intelligence" in the sense of human intelligence – there is no creativity, innovation, intuition, independent problem-solving, sentience, self-awareness, empathy, or genius. Today, even the most advanced AI systems only have the ability to analyze data and apply that analysis to solving discrete problems or to improve its capability by using its output to better define how it processes input. Whether genuinely "intelligent" AI – which moves beyond today's highly developed pattern recognition capability – will ever exist is a subject of much debate. Suffice to say that AGI doesn't exist yet. Robots aren't ready to rule the world when even the most sophisticated AI program doesn't know whether the sun rises because a rooster crows or the rooster awakens because the sun has cast its first shadows.

The breadth of the definition of narrow AI is something of a moving target. It will continue to evolve as technological change accelerates. Two decades ago, we could only marvel when IBM's Deep Blue beat world chess champion Gary Kasparov and then later when IBM's Watson defeated Ken Jennings at "Jeopardy!" We now carry in our pockets smart phones with more processing power and decision-making capabilities than that which put a man on the moon. More recently, with no human input beyond giving it the basic rules of chess and four hours of practice, Google's AlphaZero played 100 chess matches against the world's best chess playing program without losing a single match, winning 28 of them. Think of the ramifications – in the course of an afternoon, a machine moved from complete ignorance to total mastery of a game that for centuries has been considered one of the defining marvels of human intelligence. What next?

Things such as speech recognition and handwriting interpretation, seen as extraordinary a few years ago, are now viewed as mainstream technology services. This evolution will continue, as computers grow more powerful and algorithms grow more sophisticated. Experimental and exciting ideas will continue to evolve into mainstream tools.

Consider for a moment how much is changing, and how fast. The standard user interface (keyboard, mouse, screen) that we have been using for decades is increasingly being replaced by audio and video interfaces, AI at both the device level and in the cloud, and sensors that attempt to infer user intent by utilizing a variety of sensors and inference engines. Moreover, more devices will soon be processing more data – cameras, microphones, and other devices with sensors – and will be running on 5G mobile networks (ones with the capability to download a full-length motion picture in less than five seconds). Users increasingly will interact with these new devices and services through gesture, voice and human-like expressions of sight, sound and touch (and yes, smell, though that is further down the road!). As the Internet of Things (IoT) continues to grow, exponential amounts of data will be created, consumed and available for organizations to process. The cumulative interconnectedness of the IoT, more than whether your toaster can be said to "think," will raise AI concerns and opportunities.

It will not stop there. By exploiting newer augmented reality and virtual reality technologies, the ability to synthetically create complex environments and allow for human interaction will blur the lines between realities and will increasingly open a huge new set of ethical and legal challenges regarding their proper use. When combined with increasingly powerful AI engines that can quickly "learn" as these environments are explored by humans, the possibilities become endless. As do legal risks.

Consider for a moment a single data point to put in perspective the computing power that will soon be able to perform AI functions: The US Department of Energy recently put to work Summit, the world's fastest supercomputer. The race is now on between the US and China to achieve the next major milestone in computing power: a machine that would be five times faster than Summit – one that will be able to execute an exaflop, i.e., a "billion billion" calculations per second. As MIT's *Technology Review* recently put it: that amount of processing power in one machine is the equivalent of every person on earth doing a calculation every second every day for just over four years to match what that machine will be able to do in a flash.

Back to the present for a moment. Within what may currently be seen as AI using a broad frame, various functional types of it may identified:

- **Assisted Intelligence**, for example, is probably the most common form of AI. It involves computerized replacement of standardized or repetitive tasks performed by humans. This is sometimes also called "robotization," which, notwithstanding its name, does not necessarily involve machines with "limbs" performing physical tasks. One of the most common examples of robotization has been with us since the late 1960s in the form of the ATM. Assisted intelligence has the power to speed up, cheapen and make more accurate many fields of activity, but the intelligence remains primarily human.
- **Augmented Intelligence** involves humans and machines learning from each other and, through the process, redefining the boundary between which/who does what. Here, the application of learning from data processed is primarily with the machine, and the value of the breadth of data to which the machine has access is considerably greater, but a human often still sets the direction and defines the parameters. In 2010, for example, we began developing predictive coding technology to learn from lawyers which documents were (or were not) relevant to a discovery request or government subpoena. Since then, we have developed additional advanced tools and proprietary workflows to put this form of AI to work. Our Intelligent Discovery Process™ now routinely and defensibly reduces the volume of discovery review by as much as 80% as compared to traditional methods.
- **Autonomous Intelligence**, sometimes referred to as "machine learning," involves the continuous adaptive learning from analysis of data by the machine, with consequent increases in the value that can be generated from much larger and more extensive datasets. It is this that allows machines to write their own software and to play games such as Go and chess. But even in this area, where some of the most exciting recent developments have taken place, the other attributes of human intelligence are absent. These are still logic-based, data-driven systems. Thus, for example, when retail companies install machine learning call center technology (which has the capacity to deal with routine inquiries far faster and more accurately than a human call center agent), a human capability will still stand behind the technology to deal with situations that require sensitivity, such as bereavement.

We believe it is worth separating these different types or prototypes of AI because they give rise to different considerations of law and policy. As noted above, the speed of evolution of AI within the economy means that legislators, regulators, and jurists must confront highly complex questions before there is any settled view in society about the ethical framework that should underpin the role of advanced technologies. As a firm with a long history of addressing public policy issues, we recognize how future AI developments could have profound effects on business risks and opportunities, and have decades of experience in helping clients address them.

In assisted intelligence, for example, accountability today remains with the human "controller," designer, operator, installer, or programmer. For augmented intelligence, accountability also is much more likely to remain with a human, but with some interesting issues of liability between the operator and the software designer. For autonomous intelligence, these boundaries are much less clear cut. And of course, by legislation or regulatory fiat, governments can change the allocation of risk or establish liability as a matter of law. And thus it is crucial for businesses to stay abreast of these issues.

Another key issue underlying all consideration of AI is the framework for data management, which will drive enterprise value. (Although it is not a form of AI, the further development of blockchain technology also will have a profound impact on how data is managed and preserved.) AI is essentially driven by the quality of the dataset that underpins it. As AI is deployed into more complex systems (connected cars, for example, operating on smart roads), questions of ownership and ability to exploit the data produced will become increasingly acute, including from privacy and competition/antitrust perspectives, with "data monopolies" representing the new frontier of competition law, and innovative concepts such as "data trusts" providing frameworks within which data sharing can take place. The extent to which it will be possible to create genuine separation between data that could identify individuals (and which therefore would fall within privacy protection under stricter privacy regimes, such as the EU's General Data Protection Regulation) from data that could not identify individuals also will become more blurred. The question of individual identification will, in some cases, come down to how to associate different datasets, each of which could contain no personal identifiers.

Beyond these and other legal issues, society will increasingly need to address the impact of AI on skills, employment, educational opportunities and other social impacts as it is more widely deployed and has potentially more disruptive effects. As we look at that future, we believe that companies that achieve greater efficiency and value through application of AI will have to play some part in addressing those issues. And thus the need for corporate legal departments to focus on the benefits and risks of AI as it is further deployed across their enterprises with potential significant ramifications for their workers and their shareholders.

Finally, major strategic and national security developments will affect the development of AI and ultimately business risks and opportunities. For example, the Defense Advanced Research Project Agency (DARPA) recently announced a US$2 billion initiative to "explore how machines can acquire human-like communication and reasoning capabilities, with the ability to recognize new situations and environments and adapt to them." Among DARPA's goals are pioneering next-generation algorithms and applications such as "explainability" and common-sense reasoning. (The US Army and Department of Defense already have developed an algorithm to identify the ideal caffeine fix for soldiers; for the rest of us, there will soon be an app for that!) If the Defense Department is successful in endowing computers with common sense and the ability to explain the rationale for their decisions, how might that technology affect business decision making? What further ramifications might it have for society? And what are the ramifications of China setting out to be the world leader in AI by 2030? These are big issues. They will have major consequences.

We will continue to explore these consequences. We look forward in engaging in a dialogue with you as AI technology continues to reshape our world.

Contact

Jeffrey L. Turner
Partner, Washington DC
T +1 202 457 6434
E jeff.turner@squirepb.com

Matthew Kirk
International Affairs Advisor, London
T +44 207 655 1389
E matthew.kirk@squirepb.com

The contents of this update are not intended to serve as legal advice related to individual situations or as legal opinions concerning such situations, nor should they be considered a substitute for taking legal advice.

© Squire Patton Boggs.

All Rights Reserved 2018