Portfolio Media. Inc. | 111 West 19th Street, 5th floor | New York, NY 10011 | www.law360.com
Phone: +1 646 783 7100 | Fax: +1 646 783 7161 | customerservice@law360.com

How A General Counsel Should Think About AI: Part 2

Read the first part of this article here.


Bruce J. Heiman


Elana R. Reman

Today's conversations about artificial intelligence inevitably conjure up images of strong AI — of Westworld's Dolores and Ex Machina's Ava — and from there, it is but a small leap toward the ominous superintelligence. The looming question — the one that keeps people like Stephen Hawking, Nick Bostrom and Elon Musk up at night — is: when we succeed at creating a strong AI, will we still be fully in control?

This is a question that, researchers estimate, we will have decades to ponder.

The more immediate issue is what questions do you need to be asking now about AI? Now is the time to start thinking about how AI can improve corporate performance, because we have already entered the "narrow AI" revolution — augmented intelligence — where programs assist humans by performing discrete tasks faster and more accurately than a human could, from scheduling appointments, to early cancer detection, to autonomous driving. Here are five questions to get you started:

Can I Legally Repurpose My Data for AI Uses?

"The fuel for effective AI is data." AI harnesses large, ever-changing data sets and churns out findings and analytics to provide uniquely valuable insights to businesses. It teases out patterns a human could never identify. All this data poses both opportunities as well as challenges for businesses. It requires trained data scientists to identify and understand the value of an algorithm's findings, and separate out causation from correlation. But in the rush to use big data, companies must ensure that they have the rights to use that data and have complied with all applicable data privacy laws.

Data is more valuable to companies when it is personally identified, but many industries — and indeed many countries — significantly restrict its use to protect citizens against its misuse. The effectiveness of automated vehicles, for example, will be vastly increased by information sharing among all connected vehicles. No driverless car would make a mistake previously made by another driverless car, because it would learn from that mistake and adjust its situational awareness and decision-making, saving countless lives. But this information is personally identifiable to the driver or vehicle owner, and can pose great privacy risks if misused. In the U.S., where privacy is regulated sectorally, it remains to be determined what specific regulations will be passed to limit the dissemination of this information, but it is clear that privacy is a priority. In Europe, this means that "smart cars" must comply fully with the recently adopted General Data Protection Regulation, including obtaining consent from end-users, where necessary.

Likewise, mining troves of healthcare information could isolate patterns that identify hidden symptoms and causes of diseases, and hold immense promise for medical advancements to improve care coordination, estimate the costs of care, and support a myriad of public health initiatives. But the Health Insurance Portability and Accountability Act places strict guidelines on the sharing of personal health information (PHI) in the hands of HIPAA-covered entities, which include health insurance providers, healthcare providers and their business associates. HIPAA limits the types of PHI that can be collected, shared or used in marketing activities. Covered entities must "de-identify" such

information in order to share it, following HIPAA de-identification guidelines. For example, a Harvard research team used de-identified health information from a Boston healthcare system to identify adverse health events associated with diabetes medications.

Many kinds of information a company already collects in the ordinary course of business can be very useful to a company when analyzed for patterns and insights using AI software. The challenge for a general counsel is to ensure that your company can monetize the information in a way that complies with applicable privacy laws.

Can I use AI for Employment Decisions?

Data analytics was one of the top 5 human resources trends of 2016, according to Forbes magazine. Predictive analytics algorithms can help companies review applications faster than humans, identify indicators of success, increase diversity and improve retention. AI screening methods have the added bonus of eliminating the all-too-common "like me" bias of recruiters, which favors applicants who went to similar schools, have similar upbringings, etc. as the recruiter. Such biases — which AI screening promises to eliminate — lead to the hiring of like-minded individuals, increasing groupthink and reducing workplace diversity.

But, companies using artificial intelligence algorithms to assist with hiring must ensure that the results of their algorithms comply with the Equal Employment Opportunity Commission's anti-discrimination laws — including no disparate treatment or disparate impact. This can prove to be a challenge. AI algorithms pair large and evolving data sets with machine learning capabilities to analyze millions of data points of input and amalgamate them into indicators of success or failure. If the data points are, themselves, a result of human bias, the algorithm's results will only magnify that bias, unless bias-regulating mechanisms have been built in. If they go unchecked, AI algorithms will demonstrate strong confirmation bias. If most past successful business consultants were white men in their mid-thirties, a model will codify these as indicators of success. Likewise, if people with certain disabilities had high levels of absenteeism, it may view applicants with similar characteristics as bad investments for the company.

Although AI holds immense potential to improve HR practices, general counsels advising corporate HR strategies regarding the use of AI must still remember to evaluate all hiring and firing decisions under traditional anti-discrimination law.

Can I Protect AI IP?

Patentability requires that inventors be "individuals," and the Federal Circuit has held that this refers to people. Inventing also implies action, rather than a mere ownership status. Some AI algorithms respond directly to human queries and are preprogrammed to preform specific operations. Here, attributing the invention to the owner or programmer does not seem so far-fetched.

However, another type of AI applications, commonly referred to as Level D, have the ability to identify weaknesses and reprogram themselves. Level D applications use machine learning to manifest a kind of "self-awareness," and can use data they find in unique ways. With Level D applications, an eventual output is not necessarily a foreseeable result of inputs by the programmer. In these cases, the inventing seems to be done by the computer and not the person.

A question then arises whether under the current system this could be patentable. Patents are awarded to provide inventors with incentives, and some might argue that a machine does not need to be incentivized to invent. While this is true, the companies that fund AI research and development do need incentives to undertake costly research.

Another question is whether AI can infringe an existing patent? And if so, what is the remedy? These same Level D applications present a unique challenge for infringement litigation as well. Current laws governing IP protection contemplate only human infringers. A human "smoking gun" is a prerequisite for a finding of infringement. If a program such as a spider infringes copyrighted content, liability is attributed to the owner, the programmer or both, and remedies can only be sought against those people. By contrast, infringement by a Level D application cannot be traced back directly to a human, nor could a human have reasonably foreseen the infringement. The law has yet to draw a line distinguishing situations that indicate an AI app acted independently enough that the programmer

cannot be held liable. Additionally, remedies are defined in terms of monetary damages, injunctions or prison time, most of which are ineffective against an infringing machine.

Issues have already arisen with AI systems infringing patents and copyrights in their collection of mass amounts of online data. Programs that train natural language apps on collections of literary works are the clearest example. Will such uses of copyrighted materials be considered "fair uses," or will we need a "robot exception" when the use is substantially profit-motivated? Luckily, many large companies working in AI development have stemmed this issue to some extent by making many developer kits and data sets free and open. But with protected works, an AI-oriented licensing system may eventually emerge. What should the remedy be when an AI app sweeps in three protected works into a dataset of millions, and uses that data set to invent or create? These are issues on which the law is currently silent.

Can I Be Sued for Using AI?

Artificial intelligence does not fall neatly into any established liability model — and the liability model we choose will predetermine AI's future. Too strict, and we disincentivize research and development; too loose, and serious injuries go uncompensated. While some argue that Congress should step in and pass a statute to establish AI liability regimes and reduce market uncertainty, courts have historically proven capable of adapting established liability standards to emerging technologies as controversies have arisen.

Importantly, we should not allow questions about who will incur liability if something goes wrong to hold up the deployment of life-saving AI technologies that are ready for the market. Waiting until a legal standard is adopted could itself cost the number of lives that an AI technology has the potential to save. With automated vehicles, for example, that number is estimated to be about 33,300 lives per year. On average in the U.S., over 37,000 people die in car accidents every year, and of those, 90 percent are caused by human error. Automated vehicles hold the promise of eliminating these crashes, thereby saving the lives of about 91 people per day.

Determining how to assign liability in incidents involving AI technologies is an exercise in analogy. One possibility is to hold the creators or programmers of AI algorithms strictly liable for misuses and infringements caused by their algorithms. But AI can develop in a way that performs functions not foreseeable by its programmers, and holding programmers to this standard would significantly deter innovation. Defendants are usually held strictly liable in situations involving inherently high-risk activities or a significant amount of control over the entity causing harm. Examples include injuries resulting from abnormally dangerous activities, torts by employees committed within the scope of employment, and defectively manufactured products. An AI system, however, is not inherently high-risk. Likewise, its programmer can quickly lose direct control over its actions and choices.

Another possibility is to use a negligence analysis. The owners or programmers of AI applications could be held liable in those cases where they knew or should have known. The creator of a knowingly infringing algorithm would be like the keeper of a wild animal as a pet — she is strictly liable. On the other hand, the creator of a benevolent algorithm is put on notice by its first infringement of a certain type — just like the owner of a peaceful dog "who gets one free bite."

There is also the question of whose standard of care do we apply? This question is most tangible in the context of automated vehicles, where people are concerned with the tough choices algorithms will make when faced with accident conditions, or choices such as whether to avoid hitting a child or a cat. A human driver in such a situation would be held only to the standard of care of a reasonable driver faced with a split-second decision under pressure. Does an automated vehicle receive the same benefit, or is a vehicle that makes retrospectively the wrong choice under these conditions strictly liable as a defective product? If automated vehicles simply replace human drivers, is it fair to hold them to different liability standards? This question is essential to manufacturers, because building an automated vehicle that is better and safer than a human driver is easy, but building one that is perfect is nearly impossible. Perhaps a "reasonable automated vehicle" standard may eventually emerge based on the programming and ethical conventions established by the automated vehicle manufacturing industry.

Can I get AI Insurance?

The market for insurance will see significant changes with the widespread adoption of artificial intelligence technologies. Insurance is based on risk or uncertainty, and which entity can most efficiently absorb that risk. AI has the potential to vastly improve actuarial markets, which assess risk through statistical analyses of accident data, through the use of data analytics and predictive algorithms. The promise of AI for insurance is that it can more accurately predict individualized risk, thereby reducing the uncertainty absorbed by an insurance company in any given risk pool, allowing it to tie premiums more closely to actual, individual risk. While AI can greatly reduce uncertainty in some areas, it can increase it in others. Level D applications, which can learn and evolve themselves in ways unpredictable to the original programmer, can increase uncertainty because it is not yet possible to predict or control their eventual outcomes.

The risks posed by such inventions raise questions regarding whether a market for general AI insurance will eventually emerge, and what it may look like. Will current products liability insurance coverage provide the model? Many current products liability policies require the product to be finished and to require no further contribution from the manufacturer. Systems using artificial intelligence and deep learning are constantly changing and evolving, so will the products be treated as "completed" for insurance coverage purposes, or will manufacturers continue to bear risk throughout the product's entire lifecycle? Another possibility is that cyberpolicies will eventually come to include artificial intelligence and its risks as possible "cyber events."

Who ought to bear the burden when a Level D application commits a crime or causes an injury that cannot be traced back to any human input? These injuries will, no doubt, require compensation, but from whose pockets? Perhaps companies will emerge that insure the truly unpredictable, or perhaps courts will consider certain injuries "acts of god." What kinds of injuries would AI insurance cover, and how attenuated from the algorithm failure can the injuries be?

Conclusion

The artificial intelligence revolution holds great promise to make businesses more profitable and efficient. The uses of AI are virtually limitless, as are the novel legal questions it raises. General counsels face the challenging task of understanding how their companies can strategically and legally utilize their data, and navigating the web of emerging legal and regulatory rules surrounding the uses of AI. Human resources, intellectual property, liability and insurance are only some of the areas where AI will be a game-changer — in terms of both business opportunities and legal issues. General counsels must ask the right questions, and get smart on artificial intelligence.

—By Bruce J. Heiman and Elana R. Reman, K&L Gates LLP

Bruce Heiman is a partner in K&L Gates' office in Washington, D.C., and co-chairs the firm's policy and regulatory practice area. Elana Reman is an associate in the firm's public policy and law group. The firm will be conducting a series of conferences on AI during the spring, the first one of which will be in Washington, D.C., on Feb. 22, 2017.

The opinions expressed are those of the author(s) and do not necessarily reflect the views of the firm, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.