

“Deep Learning: How Intelligent Companies and Firms Challenge Artificial Intelligence”

Speakers:

James M. Lee, CEO, LegalMation

Patrick V. DiDomenico, CKO, Ogletree Deakins,

Brandon Mack, Director Analytics and Advanced Technologies, EPIQ

Moderator:

Marisa A. Trasatti, Partner, Wilson Elser Moskowitz Edelman & Dicker LLP

Federation of Defense Corporate Counsel
March 26, 2019

Artificial intelligence: What is It and What We Should Know

I. Introduction

Artificial Intelligence has often been leading edge in what has been called the Fourth Industrial Revolution. Commentators have described the varying roles of AI and the potential for disruption in each industry in which it is used. And although the notion of AI changing the world has been considered abstractly in philosophy and entertainment, federal and state policies, regulations, and laws has offered little guidance on AI.¹ Stanford's 2015 report on their One Hundred Year Study on Artificial Intelligence concludes that AI's impact by 2030 will be seen mostly in transportation, service robots, healthcare, education, low-resource communities, public safety and security, employment and workplace, and entertainment.²

With the amount of AI already in the market, the expected exponential pace of development—and the novel questions presented by the use of learning machines performing tasks previously unique to humans—understanding and planning for the legal aspects of AI is of paramount importance.³ We will provide an overview of AI, its current applications, the new laws and regulations relevant to these AI systems, and highlight the different legal and ethical considerations that may be relevant.

II. Definitions

¹ Matthew U. Scherer, Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies, 29 Harv. J.L. & Tech. 353, 357 (2016) (West).

² *Artificial Intelligence and Life in 2030, One Hundred Year Study on Artificial Intelligence (AI100)*, Report of the 2015 Study Panel, Stanford University (Sept. 2016).

³ Id.

Generally, AI is the study of “cognitive process using the conceptual frameworks and tools of computer science.”⁴ Others define AI as “that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.”⁵

Recent applications of AI can be lumped into two general categories. “Mobile AI,” like Autonomous Vehicles,⁶ or “static AI,” such as computer programs performing legal services.⁷ Another distinction is between “narrow AI”—or specific application areas such as strategic game playing, autonomous cars, image recognition and language translation—and “general AI”—which is defined as “intelligent behavior at least as advanced as a person across the full range of cognitive tasks.”⁸ Recent developments in AI applications have been in narrow AI, whereas experts believe progress on general AI is still at least decades away.⁹

III. How does AI Work?

a. Pattern recognition

Machine learning techniques, where AI systems learn by examples, or teach themselves, to carry out tasks based on pattern recognition, are broken down into four categories.¹⁰ Deep learning is the “use of software algorithms to analyze large datasets in what are called neural networks as they seek to mimic the way the human brain works.”

⁴ Edwina L. Rissland, Artificial Intelligence and Law: Stepping Stones to A Model of Legal Reasoning, 99 Yale L.J. 1957, 1958 (1990) (West).

⁵ Nils J. Nilsson, The Quest for Artificial Intelligence: A History of Ideas and Achievements (Cambridge, UK: Cambridge University Press, 2010).

⁶ See generally Jessica S. Brodsky, Autonomous Vehicle Regulation: How an Uncertain Legal Landscape May Hit the Brakes on Self-Driving Cars, 31 Berkeley Tech. L.J. 851 (2016) (West).

⁷ Artificial Intelligence and Life in 2030, *supra* note 2 at 14.

⁸ Preparing for the Future of Artificial Intelligence, Executive Office of the President National Science and Technology Council Committee on Technology, at 7 (October 2016).

⁹ Id.

¹⁰ Kemp, *supra* note 13 at 3.

Supervised learning is the method of training an AI system “with a restricted dataset of labelled examples.”

Unsupervised learning is when AI systems are exposed to huge volumes of unlabeled data, and allows the system to make up its own rules for what to look for, therefore allowing the AI to discover otherwise hidden correlations in data.

Reinforcement learning is a combination of supervised and unsupervised learning where machine learning starts training by examples and reduced datasets, and then allows the AI to learn by itself with unlabeled data.

Lastly, large-scale machine learning is about scaling machine learning by introducing algorithms to larger datasets “so that algorithms need only operate once (as opposed to several times) on all, or even part of, the data to achieve a faster response.”¹¹

b. Human element / involvement

Reports about AI systems reflecting human biases, such as Google’s “word2vec” linguistic reconstruction algorithm reflecting gendered biases,¹² or Microsoft’s failed Twitter-based chatbot, “Tay”¹³ illustrate that the process of machine learning to an extent will always be affected by the “human element” of our biases in certain data used to teach AI.¹⁴ On the other hand, ominous reports about Facebook AI Research Lab (FAIR) shutting down an AI system after bots developed their own language to communicate—which may have been exaggerated¹⁵—highlight the loss of the human element of control in unsupervised machine

¹¹ Id.

¹² Levendowski, *supra* note 55.

¹³ Newman, *supra* note 54.

¹⁴ Levendowski, *supra* note 55.

¹⁵ Tom McKay, *No, Facebook Did Not Panic and Shut Down an AI Program That Was Getting Dangerously Smart*, (Jul. 2017) <https://gizmodo.com/no-facebook-did-not-panic-and-shut-down-an-ai-program-1797414922>

learning. Recognizing the challenge of eliminating data bias, and the unknowns of unsupervised machine learning, human involvement—from programmers to professionals supervising the development and operation of AI systems—is both inevitable and necessary to make better, unbiased AI. Expanding access to better data for supervised and unsupervised machine learning is one of the key areas of human involvement.

c. Corpus concept

In AI, the “corpus” refers to the datasets being used to teach an algorithm through machine learning.¹⁶ In the case of Microsoft’s twitter bot experiment, Tay, the corpus of information used for training was human created works in Google News.¹⁷ That information corpus has been critiqued as reflecting the sexist biases of the human authors of the works available in Google News.¹⁸

A recent law review article argues that one way to allow for expanded access to better data for machine learning—to reduce data biases—is through the copyright doctrine of fair use.¹⁹ Using more copyrighted works to train AI systems, though not unbiased themselves, has the advantage of increasing the total amount of datasets, which allows for more accurate performance. It also allows researchers and developers to supplement a corpus with new data to balance out any identified biases, in addition to other advantages.²⁰ In the development of a corpus, then, there is a clear tension between balancing the interests of the authors of copyrighted content and the interests of those developing and creating AI systems, as well as the public.²¹

¹⁶ Id.

¹⁷ Levendowski, *supra* note 55.

¹⁸ Id.

¹⁹ Id.

²⁰ Id.

²¹ Id.

IV. Legal Considerations

Virtually every industry is already being affected by AI in some fashion;²² common concerns for developers include mitigating data bias to build and maintain public confidence in AI, as well as data privacy and cybersecurity.²³ Access to a sufficient amount and quality of data to use for training AI systems is also relevant for developers and operators of AI systems trying to mitigate biases. Current copyright law appears to favor holders of works suitable for use as training data, and thus big AI investors who own platforms supplying them with data, over other developers having to negotiate for access to data. Although untested arguments have been made that the transformative purpose of human created works being used as merely raw data should already qualify as fair use.²⁴ Given the amounts of vast amount of personal data generated and shared in the development of AI systems—from Autonomous Vehicles knowing your daily travels, to vacuums mapping your entire house—privacy and cybersecurity measures are also shared concerns among different AI.

a. Data bias

AI is not perfect. Developing and improving AI requires training data. The biases in data used to train AI can result in false positives and false negatives, leading to decreased public confidence in the positive potential of AI.²⁵ Evidence includes public backlash against sexist chatbots,²⁶ facial recognition algorithms misidentifying people,²⁷ and complaints that privately

²² Littler Mendelson's Workplace Policy Institute, Prime Policy Group, *Comments on the Department of Labor's Draft Fiscal Years 2018-2022 Strategic Plan: Incorporating the Impact of Artificial Intelligence, Robotics, and Other Automated Systems Technologies into DOL's Strategic Goals*, (Dec. 2017) <http://prime-policy.com/wp-content/uploads/2017/12/DOL-Comments-re-Automated-Systems-Final.pdf>

²³ See *supra* chapter I section c.

²⁴ See generally Levendowski, *supra* note 55, (analogizing to other disruptive technologies to copyright which led to fair uses such as time-shifting).

²⁵ Issie Lapowsky, *One State's Bail Reform Exposes the Promise and Pitfalls of Tech-driven Justice*, Wired.com (Sept. 2017) <https://www.wired.com/story/bail-reform-tech-justice>

²⁶ Newman, *supra* note 54.

developed risk score algorithms for criminal defendants licensed for use by law enforcement agencies are racist.²⁸ Developers of AI systems should to be conscious of these issues by expanding and adapting their corpus of training data accordingly to mitigate identified biases. Some developers may need to negotiate contract and license agreements for the use of training data—from nonfictional and fictional works, to Facebook selfies used for facial recognition—protected by copyright. Also, records of data used to combat bias is another consideration, as the use of IP protected works as training data presents unresolved liability questions. Internal policies should be developed regarding whether, or the extent to which, an AI system’s algorithm and underlying training data will be disclosed publicly.²⁹ Disclosure of the algorithm may foreclose potential intellectual property protection, and disclosure of the underlying training data raises unanswered copyright liability questions.

b. Poor training / supervised training

Modern AI systems based on data analytics are taught to think with books, articles, photographs, films and audio recordings available on the internet.³⁰ There is evidence that using low risk sources of training data such as Wikipedia or outdated public domain works results in demonstrable biases in the performance of AI systems.³¹ Developers should consider the need to expand the corpus of training data to include more contemporary “data”—books, articles, photographs, films and audio recordings—relevant to a given AI project. Expanding a corpus to address data bias implicates using works protected by intellectual property, namely copyright.

²⁷ Levendowski, *supra* note 55.

²⁸ See generally Lapowsky, *supra* note 78; Levendowski, *supra* note 55.

²⁹ See Levendowski, *supra* note 55 (Among the big data and AI players, disclosing their algorithms is more common than the underlying training data; Google, IBM, and Microsoft have released open source algorithms but not the underlying training data).

³⁰ Id.

³¹ Id.

With potential statutory damages meaning copies of protected works used to train AI could potentially result in \$150,000 in damages per copy, those developers not owning a platform flooding them with data such as Facebook, have to tread carefully. To date, no court has decided whether a copy of a protected work used to train AI qualifies as a “copy” under the 1976 Act, or more importantly, whether that would violate the right of reproduction and constitute infringement.³² Until those questions answered by a court, or new legislation,³³ developers with access issues wishing to correct identified data biases will have to rely on contract and licensing agreements for the use of copyright protected data, or continue the potentially untenable trend of blackboxing, or not releasing the algorithms and the underlying data used to train AI.³⁴

c. Case law overview

AI systems are already performing a host of complex tasks without human control, and the scope of those tasks is only increasing.³⁵ Today, AI technology can build your investment portfolio,³⁶ diagnose a disease,³⁷ and even handle early stages of litigation including initial document drafting.³⁸ The various state definitions of AI as applied to vehicles, and the seemingly imminent federal regulations, are a welcome signal of the legal system adapting to AI. Some cases may have already planted the first guideposts, however.

Regarding legal services, a Second Circuit case interpreting the meaning of “engaged in the practice of law” in a regulation exempting lawyers from mandatory overtime pay under the

³² Id.

³³ See id. (arguing for the applicability of fair use to copies of protected works used for training data, and for creating a “fair use” exception in the Computer Fraud and Abuse Act).

³⁴ Id.

³⁵ Scherer, *supra* note 1 at 363.

³⁶ Id.

³⁷ Rissland, *supra* note 3 at 1958.

³⁸ *See* Miranda Katz, *Welcome to the Era of the AI Coworker*, Wired.com (Nov. 2017), available at <https://www.wired.com/story/welcome-to-the-era-of-the-ai-coworker/>; *see also* Ben Hancock, *James Lee: Artificial Intelligence and the Future of Litigation*, Law.com (Oct. 2017), available at <https://www.law.com/sites/almstaff/2017/10/02/james-lee-artificial-intelligence-and-the-future-of-litigation/>.

Fair Labor Standards Act spawned what has been referred to as the “Lola Machine Test.”³⁹ The Court held that the plaintiff sufficiently alleged that the document review he contracted to perform did not require any exercise of legal judgment, and was therefore not “the practice of law.”⁴⁰ The test arose from the court highlighting each party’s acknowledgment at oral argument that reviewing discovery documents, which could otherwise be accomplished entirely by a machine, cannot be considered the practice of law. This principle has implications for firms incorporating AI systems such as ROSS and LegalMation.

A federal district court ordering consolidation of various putative class actions held that the use of AI technology in a website—to mimic interactions with real people to users of the site—could result in liability for fraud.⁴¹ A tort case involving a plane crash that occurred while the plane was on autopilot held the designer and manufacturer of the aircraft—as opposed to the designer and manufacturer of the autopilot system—liable for a manufacturing defect.⁴² Other areas where the legal system has interacted superficially with AI include claims by workers injured on the job by robots, and about the safety of surgical robots.⁴³

In 1984, the Third Circuit Court of Appeals held that “robots cannot be sued.”⁴⁴ Due to the cost advantages of using open source and Commercial off the shelf (COTS) software and hardware, however—and the reality therefore that not all “creators” or “operators” of AI technology will have written entirely bespoke code (or created every component) for their AI—

³⁹ David Horrigan, *The Lola Machine Test: Should Courts Let Technology Define the Practice of Law?*, *relativity.com*, (August 2015) <https://www.relativity.com/blog/the-lola-machine-test-should-courts-let-technology-define-the-practice-of-law/>

⁴⁰ *Lola v. Skadden, Arps, Slate, Meagher & Flom LLP*, 620 Fed. Appx. 37, 45 (2d Cir. 2015) (unpublished) (West).

⁴¹ See *In re Ashley Madison Customer Data Sec. Breach Litig.*, 148 F. Supp. 3d 1378, 1380 (U.S. Jud. Pan. Mult. Lit. 2015) (West) (plaintiffs alleging AI software responding to male user inquiries as female users was fraudulent).

⁴² See *Ferguson v. Bombardier Services Corp.*, 244 Fed. Appx. 944, 951 (11th Cir. 2007) (unpublished) (West) (significantly, however liability in this case was imposed on the designer / manufacturer of the aircraft because they actually installed the AI system on the airplane).

⁴³ James S. Azadian & Garrett M. Fahy, *Artificial Intelligence and the Law: Navigating "Known Unknowns"*, 59 *Orange County Law*. 22, 29 (October 2017) (West).

⁴⁴ *U.S. v. Athlone Industries, Inc.*, 746 F.2d 977, 979 (3d Cir. 1984) (West).

regulating and imposing liability when new AI systems are a mish mash of components from unrelated parties presents significant problems of proof and causation.⁴⁵ While similar issues have been resolved by the legal system with respect to products such as cars,⁴⁶ which also include numerous products from multiple companies, modern learning AI technologies present harder causation issues requiring value judgments the law will have to make.⁴⁷

d. Legislative attempts

Legislation at the federal and state levels has been outpaced by AI research and development. Therefore, the terms of private agreements between developers, operators and end users will likely govern much of the AI systems in the market. Some signals, like the House of Representatives passing of the SELF-DRIVE Act, show federal legislation in the arena of Autonomous Vehicles may be coming. Currently, the majority of AI projects do not have significant issues with regulatory compliance—save for Autonomous Vehicles and medical AI, which have unique barriers to innovation—likely due to the rapid development and amorphous nature of AI applications.⁴⁸ Some experts are concerned that the current administration’s policies are not conducive to the development of AI, or workforce preparation for the different kinds of jobs servicing and maintaining AI systems will demand.⁴⁹

V. Specific Industries

a. Automotive

⁴⁵ Scherer, *supra* note 1 at 371.

⁴⁶ *See Id.* at 374 (discussing multiple tortfeasor liability for indivisible harm, and contribution and indemnity).

⁴⁷ *See Id.* at 366-367 (asking whether potential harm caused by learning-AI systems acting on acquired knowledge should be considered unforeseeable, and an intervening cause absolving the system’s designers of liability).

⁴⁸ Oren Etzioni, *How to Regulate Artificial Intelligence*, nytimes.com (Sept. 2017)

<https://www.nytimes.com/2017/09/01/opinion/artificial-intelligence-regulations-rules.html>

⁴⁹ *See generally supra* note 75.

Today's semi-autonomous vehicle systems, and "cobot" drones and robots used in warehouses, are expected to become tomorrow's fully autonomous intelligent transportation systems, and online-purchase revolutionizing delivery drones.⁵⁰ A growing minority of states, at least nine states to date, have laws or regulations addressing Autonomous Vehicles.⁵¹ Some, like Michigan, Tennessee and Georgia have passed general authorization statutes for the testing of Autonomous Vehicles if they otherwise comply with applicable vehicle laws and operators secure proof of sufficient insurance.⁵² Others, like California, have more extensive regulations in place which offer some of the first statutory or regulatory definitions of AI systems.⁵³ California also appears to be the first state to authorize in one county the testing of fully Autonomous Vehicles on public roads, meaning vehicles without any driver present to take fallback control.⁵⁴

Unless Congress is ultimately able to come together and pass an Autonomous Vehicles bill, which some are skeptical of given the differences in the Senate's draft bill,⁵⁵ regulation of Autonomous Vehicles will remain at the state level.⁵⁶ The biggest concerns for developers of AV's then will likely remain compliance with the different, and sometimes conflicting state vehicle laws, as well as shared issues relating to privacy, and cybersecurity.⁵⁷ If the House bill provides any indication, federal legislation would give the NHSTA the exclusive regulatory power over the design, construction, and performance of Autonomous Vehicles, as it currently

⁵⁰ See generally *supra* note 2.

⁵¹ See Cal. Veh. Code § 38750; C.G.S.A. P.A. 17-69, § 1 [Connecticut]; D.C. Code Ann. § 50-2352; Fla. Stat. Ann. § 316.85; Ga. Code Ann. § 40-6-279; Mich. Comp. Laws Ann. § 257.665; Nev. Rev. Stat. Ann. § 482A.030; Tenn. Code Ann. § 55-8-202; Utah Code Ann. § 41-26-102 (West).

⁵² See generally *id.*

⁵³ See *supra* chapter I section III(a).

⁵⁴ Veh. Code. Contra Costa County § 38755.

⁵⁵ Aarian Marshall, *Congress Unites (Gasp) to Spread Self-Driving Cars Across America*, wired.com (Sept. 2017) <https://www.wired.com/story/congress-self-driving-car-law-bill/>

⁵⁶ See *supra* chapter I section III(a).

⁵⁷ Brodsky, *supra* note 8 at 871.

has with normal cars, leaving states to regulate registration, and licensing.⁵⁸ The House bill also would mandate that manufacturers maintain detailed written privacy plans regarding the collection, use, sharing and storage of user data, giving express jurisdiction to the Federal Trade Commission to regulate violations as an “unfair or deceptive practice” under § 5(a)(1) of the FTC Act.⁵⁹ A similar requirement that manufacturers maintain cybersecurity plans is also included.⁶⁰ Lastly, the substantial increase in NHSTA exemptions proposed under the House bill outlined above⁶¹ illustrates a potential federal policy of flexibility for increased testing, while ensuring substantial data protection and cybersecurity measures are in place.

b. Law

The legal services industry in the United States is approximately a \$275 billion market annually.⁶² AI systems performing early stage litigation services, like ROSS and LegalMation, are crossing the nonbinding boundary offered by the Second Circuit’s *lola* machine test for determining what qualifies as “the practice of law.” The host of ethical rules governing lawyers, combined with important questions that remain unanswered as courts are slowly acknowledging that the practice of law is changing, reinforces that the incorporation of legal service AI systems should be carefully considered. Designers and operators should establish policies and procedures to navigate the ABA Model Rules of Professional Conduct and relevant state bar rules, in order to avoid sanctions from the use and training of legal service AI systems. Professionals operating

⁵⁸ *See supra* note 27.

⁵⁹ 15 U.S.C 45(a)(1).

⁶⁰ *Id.*

⁶¹ *See supra* note 27.

⁶² James E. Daily, Embracing New (and Old) Ideas, 53 Wash. U. J.L. & Policy 157, 162 (2017).

these systems, lawyers, should ensure they have a prominent voice in the development of these AI systems to achieve this end.⁶³

Considerations in the legal industry include ensuring sufficient client communication about a firm's use of AI,⁶⁴ that reasonable steps are taken to maintain effective measures assuring all lawyers in a firm, including those operating AI systems and their supervising attorneys conform to the Model Rules—particularly regarding confidentiality, and duties to former clients.⁶⁵ Other likely ethical issues with the Model Rules presented by the use of AI include the interplay between the restrictions on non-lawyer assistance, the professional independence of lawyers and the need to exercise independent judgment and give candid advice. While concrete guidance to each of these concerns is not yet available, reasonable policies and procedures can be developed nonetheless. For example, the responsibilities of partners and supervising attorneys will likely be expanded to include responsibility for AI drafted briefs, document review, and so on.

Deciding whether to restrict the use of legal service AI systems to legal practitioners or allow general public access presents even more ethical issues. Should they choose to, designers and developers should consider employing lawyers to answer inevitable consumer questions, and ensure the quality of AI work products, instead of relying on non-lawyers to answer questions which may blur the line on giving legal advice, and thus raise ethical issues about the unauthorized practice of law.⁶⁶

Lastly, lawyers developing, supervising and operating AI systems should be conscious of biases—to ensure the competent performance of AI systems—and address confidentiality

⁶³ See *supra* note 75.

⁶⁴ MRCP Rule 1.4.

⁶⁵ MRCP 1.6; 1.9.

⁶⁶ See Scott B. Garner, Artificial Intelligence and Its Not-So-Artificial Legal Ethics Implications, 59 Orange County Law. 64, 66 (October 2017) (discussing the ethical issues surrounding legalzoom).

concerns raised by the use of client data.⁶⁷ Cases, briefs, and other training data should be considered for not only their legal validity, but other externalities like social consciousness, to avoid skewed data and potentially fatal mistakes in legal analysis.

c. Medicine

Healthcare and medical applications of AI systems are expected to improve quality of life for millions.⁶⁸ AI systems in medicine are already assisting with “clinical decision support, patient monitoring and coaching, automated devices to assist in surgery or patient care, and management of healthcare systems.”⁶⁹ Legal considerations about the use of static AI and mobile, robotic AI in medicine, begin with the reality that development and deployment of these systems has been hamstrung by outdated regulations and slow approval by the relevant regulatory agencies.⁷⁰ New innovations in diagnostic software algorithms have received slow to approval by the FDA, for example, but other legal considerations present barriers to development for AI in this area as well. Continuing the theme of the central importance of data for modern AI systems—the Health Insurance Portability and Accountability Act (HIPAA) patient privacy requirements present obstacles to the accessibility and flow of patient data. To use the example of wearable diagnostic software algorithms, identifying potentially harmful drug interactions may be prevented due to blocked access to patient records.

In the domain of medicine and healthcare, then, legal considerations for AI developers and operators continue to be access to a sufficient amount of quality training data, protecting the

⁶⁷ See *supra* chapter II

⁶⁸ See *generally supra* note 2 at 25.

⁶⁹ *Id.*

⁷⁰ *Id.* at 26

privacy of all user information collected, and cybersecurity. Regulatory reform may be the first important hurdle, however.⁷¹

d. Agriculture and food preparation

Other industries beginning to see the influence of AI and robotics include industrial agriculture, food preparation and food services. Whereas AI in areas like professional services are feared for the prospect of taking away jobs—though some experts stress that AI will only change the nature of the jobs in demand⁷²—AI applications in agriculture are expected to correct for the problem of not attracting young workers. On the other hand, the AI threat to minimum wage jobs in food preparation is well documented.⁷³ AI systems in agriculture and food preparation are predominantly mobile AI incorporating robotic hardware, and are beginning to work alongside human employees. Injuries caused by robots on the job have already led to lawsuits.⁷⁴ This raises unsettled legal questions about whether AI acting on acquired knowledge through unsupervised learning should be considered an unforeseeable, intervening cause, potentially relieving a designer or operator of tort liability. Current, and outdated OSHA guidelines are also highlighted by experts as raising potential legal issues for the adoption of AI bots in agriculture and food services.⁷⁵ As with every industry adopting AI, cybersecurity measures also will play a central role. Thus, allocating risk between developers, operators and employees working with these AI systems should be addressed to the extent it can be in private agreements, without developed case law on point.

e. Government uses of AI

⁷¹ See generally *id.*

⁷² See generally *supra* note 75 (arguing that demand for servicing AI jobs and cybersecurity workers will rise).

⁷³ *How Food-Bots Are Changing How We Eat*, wired.com, <https://www.youtube.com/watch?v=SKBHnbYo-4s>

⁷⁴ *Supra* note 46.

⁷⁵ See generally *supra* note 75.

The increasing use of AI systems by public agencies is perhaps the most controversial “industry” application. The need for building public trust and confidence in AI systems, while shared among all AI, is particularly relevant in this context. From criminal justice, to health and welfare, AI is being used in scoring systems and algorithms to inform decision makers on life-changing decisions, such as granting bail, sentencing, and prioritizing services.⁷⁶ Public criticism has been aimed at systems like Northpointe’s COMPAS and the Arnold Foundation’s “data-driven decision making” tool, which predict a criminal defendant’s risk factors for things like recidivism, and determine bail eligibility. Studies done by non-profits like ProPublica suggest that the risk assessments produced by these AI systems are racially biased, and consistently less accurate for African Americans, women, and young people.⁷⁷

The director of the ACLU’s criminal law reform project highlighted the omnipresent theme with modern AI, “[a]lgorithms and predictive tools are only as good as the data that’s fed into them.”⁷⁸ Transparency issues surrounding the algorithms and underlying training data used with these AI systems have led to lawsuits and investigative journalism inspiring public mistrust. Perhaps more important than in any other domain adopting AI, developers creating systems for use by public agencies should be conscious of balancing the need to build public trust—by increasing the extent to which they disclose their algorithm and underlying training data—against their intellectual property interests.

⁷⁶ See generally Lapowsky, note 80; and Levendowski, *supra* note 55 (giving the examples of a Boston man who sued a registry after having his license revoked when it falsely identified him, and the Taiwanese student who was stranded in an Australian airport when he couldn’t renew his passport because AI incorrectly identified his eyes as being closed).

⁷⁷ See generally Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, *Machine Bias: There’s software used across the country to predict future criminals. And its biased against blacks*, ProPublica (May 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

⁷⁸ See generally Lapowsky, note 80.