**Regulation and Ethics of Artificial Intelligence**

*Recognizing the Benefits of AI Without Falling Victim to its Avoidable and Manageable Shortcomings*

By Debra J. Hall*

## I. INTRODUCTION

"Recent technological change has transformed almost every part of life. Today, technology influences our relationships, decisions, desires and the way we experience reality."[1]

-- Matthew Beard & Simon Longstaff

Artificial intelligence (AI) will continue to revolutionize every facet of our lives and have a profound impact on the world economy,[2] including the way we think about risk and liability. But in reality, we are only witnessing the beginning days of this transformative technology. The Defense Advanced Research Projects Agency (DARPA) recently announced its focus on the "third wave" of AI theory and application which will transform computers from specialized tools into machines with "human-like communication and reasoning capabilities, with the ability to recognize new situations and environments and adapt to them."[3]

In July 2017, the Comptroller General of the United States convened a Forum on Artificial Intelligence, including participants from industry, government, academia, and non-profit organizations. Forum participants highlighted a number of challenges related to AI, including: data bias; issues relating to the collection and sharing of data needed to train AI systems; the adequacy of current laws and regulations; and the need to develop and adopt an appropriate ethical framework to govern the use of AI in research. This article addresses these topics as well as touching on the regulatory implications to the insurance industry arising from AI development and adoption.[4]

## II. AI AND MACHINE LEARNING

Many important decisions historically made by people are now made by computers. Algorithms count votes, approve loan and credit card applications, target citizens or neighborhoods for police scrutiny, prepare taxes, select taxpayers for IRS audit, grant or deny immigration visas, help identify serial rapists by reducing the turnaround time on untested rape kits, prepare patent claims and even invent new patents, aid radiologists in detecting wrist fractures and other imaging diagnostics, settle insurance claims and empower the advent of driverless cars.[5]

### *Machine Learning*

This article focuses primarily on the subset of AI known as machine learning. Modern machine learning applies and refines a series of algorithms on a large data set by optimizing iteratively as it learns in order to identify patterns and make predictions for new data.[6] Data may be of different types and qualities and may be obtained from different sources ("structured" as in an explicit database or "unstructured" e.g., information may be obtained from diverse sources on the Internet, massive amounts of pictures or other data). Computers develop these abilities either from "learning algorithms," written by humans who feed massive amounts of training data into

an artificial neural network[7] (named for its ability to process information in a way that is loosely based on the brain's nerve cell structure) or through no human intervention at all.[8]

An excellent example of both the power of AI today and the difference between supervised and unsupervised learning is DeepMind Technology's latest evolution of AlphaGo, the first computer program to defeat a world champion at the ancient Chinese game of Go. AlphaGo initially trained on thousands of human amateur and professional games to learn how to play Go. But the new version of the program, called AlphaGo Zero, skips this step and learns to play simply by playing games against itself, starting from completely random play.

After just 3 days of learning through self-play, AlphaGo Zero defeated the previous version AlphaGo (which had itself defeated the human world champion 18 times) by 100 games to 0. As noted by DeepMind Technology's CEO, "[o]ver the course of millions of AlphaGo vs AlphaGo games, the system progressively learned the game of Go from scratch, accumulating thousands of years of human knowledge during a period of just a few days."[9]

The AlphaGo Zero example is what some researchers refer to as "machine teaching" – and what commentators suggest will be the biggest exponential leap for AI – machines teaching one another. Imagine a machine that has taught itself in a number of days what humans learned over thousands of years, as with the game of Go – and then transferring that knowledge to another machine with the same potential.[10]

## III. IMPLICATIONS FOR (RE)INSURANCE PROFESSIONALS

As insurance and reinsurance professionals we are at the epicenter of both the legal and regulatory impact of AI. We can expect: new insurance products;[11] modifications to existing products;[12] new ways of underwriting products,[13] modified distribution systems;[14] new ways of pricing products;[15] changes in the way that we look at risk, to including "risk slicing;"[16] new risks,[17] and changes to current revenue streams.[18]

At the same time, we can count on regulators examining the impact and effect on consumers' privacy, insurance risk and pricing, the potential for new non-insurance players in the market, insurance solvency, and much more to drive the regulation of the use of AI by those in the industry. Whether insurers will be able to play on a level playing field with others, be stifled in their application of AI, or be scrutinized and tested in unanticipated ways is yet to be seen.

Although data plays a central role in the insurance industry, it is estimated that most insurers only process 10-15% of the data they have access to, mostly structured data housed in traditional databases.[19]And as insurers struggle with creating, monitoring and implementing AI within their own business, they will need to be ever-vigilant about the manner in which they underwrite clients going forward who are also utilizing AI in their businesses. Reinsurers will be wise to fully understand how their ceding companies are underwriting and monitoring the AI risk they insure.

Disputes between policyholders and insurers and between insurers and reinsurers will no doubt arise and need to be arbitrated or litigated. Given the proprietary and confidential nature of AI, it

may behoove these industry players, to utilize the confidentiality afforded by arbitration. By necessity, arbitrators should be aware of the AI landscape and ready to take on these challenges.

## IV. UNDERSTANDING THE MAJOR CHALLENGES OF AI

### *The Problem of Bias*

AI systems rely on huge amounts of data, making it essential to understand how the data is influencing the behavior of the AI system. For example, if the system is trained on biased data, it will make unbalanced or unfair decisions, which may favor some groups over others. The irony is that the ability of machine learning to analyze data at a very granular level has the potential to produce more accurate pricing and risk assessment but is challenged by outcomes that might implicitly correlate with the discriminatory characteristics that industry regulators seek to prohibit.[20]

At its core, data bias is how discrimination of various types is translated into technology. This is not to suggest that this is the intended outcome. Bias can be inadvertently introduced in numerous ways, including:

- lack of diverse thought or experiences among those who are training the AI;[21]
- framing of the algorithmic model (what the data scientists want the AI to achieve);[22]
- training data can be either unrepresentative of reality or can itself reflect existing prejudices;[23] or
- during the data preparation stage, bias can exist in the selection of attributes chosen for the algorithm to consider,[24] which may include use of proxies that, in effect, introduce the discriminatory factors into the AI in indirect ways.[25]

Potential bias is not limited to race or gender, but extends to the economic background of the technologists, their religious preferences, as well as the full spectrum of their own experiences.[26] Experts assert that the key to diversifying data, and thereby minimizing bias, is to diversify the human beings that are accountable for the data in the first instance, thinking critically about data to ferret out biases, owning the process and taking responsibility for the consequences.[27]

Others suggest that bias is harder to fix than simply ensuring diversity in people and data on the front end. Fixing data bias involves trying to predict and identify downstream bias impacts before it's too late; altering the standard testing practices that might mask bias in the training/validation process; avoiding the "portability trap" in which a system designed for one purpose or geographic area might not be fairly used in another; and indeed, defining what is "fair."[28]

Fortunately, some AI researchers are hard at work addressing these problems by creating algorithms to detect and mitigate biases hidden within training data or learned by the model regardless of the integrity of the data; processes that hold companies accountable for fair outcomes from their systems and discussions aimed at discerning the various definitions of "fairness."[29] Accenture recently introduced an "AI Fairness tool" which uses AI to examine how data influences variables such as age, gender and race in a given model.[30]

The Congressional report, *Rise of the Machines*, warned that as AI is increasingly deployed into industries such as finance, law and medicine, existing biases reinforced by technology can cause harm to populations. Transparency is key to identifying bias – not only with respect to the system itself and the data the algorithm relies upon, but also how and why it makes the decisions it does. Transparency in this context is sometimes referred to as "interpretability" or "explainability."[31]

***The Opacity or Black Box Problem***

Machine learning techniques have the potential to achieve a high degree of accuracy and avoid the errors that might be made by humans. However, the complexities of these AI systems, and the basis upon which they make decisions, often elude humans, including those who created the systems. Often referred to as opaque or black box AI, it is sometimes not possible to track the reason that the system made the decision it did.

Some might question why results need to be explainable? After all, driving a car does not require the knowledge to build one. The need for explainability should be examined in connection with the impact that the technology is expected to have in the real world. If the AI is used for finding songs or movies to entertain us, interpretability doesn't matter much but when there are implications affecting our health, safety, or finances—interpretability becomes very relevant.

The concern about transparency is particularly important with respect to what is known as "deep learning."[32] While deep learning has proved very powerful in recent years and the hope is that it will play an essential role in diagnosing deadly diseases and solving some of the most challenging problems that societies face – this won't happen – and shouldn't happen – unless we can make these systems more understandable to their creators and accountable to their users.[33] Humans want to know why AI made a given decision—particularly when the decision impacts a major life decision, or even life itself. The OECD notes that "millions or even billions of parameters used by deep learning to solve a problem do not easily allow its results to be reverse-engineered."[34]

The Financial Stability Board (FSB) has warned that the lack of "interpretability" or "auditability" of AI and machine learning methods could present a macro-level risk if not appropriate supervised.[35] This lack of interpretability could be even more problematic during a systemic shock. Although recognizing the scarcity of skilled resources as a problem, the FSB nonetheless recommends that beyond the staff operating the AI applications, there needs to be oversight by key functions including risk management, internal audit, administrative management and regulators.[36]

> It is important that progress in AI and machine learning applications is accompanied with further progress in the interpretation of algorithms' outputs and decisions. Increased complexities of models may strain the abilities of developers and users to fully explain, and/or, in some instances, understand how they work. Efforts to improve the interpretability of AI and machine learning may be important conditions not only for risk management, as noted above, but also for greater trust

> from the general public as well as regulators and supervisors in critical financial services.[37]

Interestingly, AI technology is being developed as a way to interpret the rationale for the decisions made by other AI systems. Mark Riedl, director of the Entertainment Intelligence Lab at the Georgia Institute of Technology in Atlanta coined the term "AI Rationalization" to describe how we can train a second parallel Neural Network to semantically describe the actions of the first. [38] Using a 1980s video game, called Frogger, Riedl asked human subjects playing the game to describe their tactics aloud in real time, then recorded those comments in the game's code, trained a second system to translate between the two from code to English, producing an AI that would translate into human terms the decisions it made about the frog's movement.

Similarly, researchers believe that AI will play a critical role in helping us defend against cyber attacks–another example of machines helping us to address the problems of other machines.

**The Challenge of Controlling Adversarial AI**

One concern for AI experts and researchers is the malicious use of AI, or "adversarial AI." Imagine AI being manipulated to read benign tumors as malignant in order to advance an insurance fraud scheme, AI changing or deleting stop signs so that autonomous vehicles crashed into each other,[39] or AI-generated "deep fake content"—the generation of text or video which could be mistaken for plausible news stories or events.[40]

In fact, deep fake content has gotten so much recent attention that U.S. Senator Ben Sasse (R-NE) introduced legislation in December 2018 to criminalize the malicious creation and distribution of deep fakes.[41] Although the bill expired at year-end, it is likely to be re-introduced. Legislation addressing deep fakes was also introduced in 2017 in the New York State Assembly and reportedly opposed by several Hollywood companies.[42] One free-market think-tank commentator has suggested that deep fakes are nothing new, citing to the positive uses of the technology and predicting that society will learn to mitigate any potential resulting harm without the need for regulation or legislation.[43]

To provide more focus and coordination to an effort to fight adversarial AI that has to this point been ad hoc, IBM Research Ireland has released the Adversarial Robustness Toolbox, an open-source software library to support both researchers and developers in defending against adversarial attacks in the hope of making AI systems more secure. IBM defines the adversarial AI threat on their website as follows:

> Adversarial attacks pose a real threat to the deployment of AI systems in security critical applications. Virtually undetectable alterations of images, video, speech, and other data have been crafted to confuse AI systems. Such alterations can be crafted even if the attacker doesn't have exact knowledge of the architecture of the [AI system] or access to its parameters. Even more worrisome, adversarial attacks can be launched in the physical world: instead of manipulating the pixels of a digital image, adversaries could evade face recognition systems by wearing specially designed glasses, or defeat visual recognition systems in autonomous vehicles by sticking patches to traffic signs. [44]

Those contracting with AI vendors and incorporating AI into their own systems should take heed—designing an AI system ethically is not enough—it must also resist *unethical* human interventions.[45]

## V. THE ETHICS OF AI

While AI holds the promise of solving some of the most intractable problems of our time – it presents unique and sometimes vexing challenges. Technologists have observed: there is no reason to think we are obliged to choose between scientific advances and ethics. But, as with any project design, you can't solve problems you don't acknowledge.[46]

The World Economic Forum has defined AI as the software engine that drives the "Fourth Industrial Revolution"[47] but has cautioned that we proceed intentionally and ethically, recognizing that decisions regarding responsible AI design are often made by engineers "with little training in the complex considerations at play."[48]

> Creating AI without considering the potential ethical and human-centered implications creates liabilities for the evolution of social, economic and governance systems. In view of the magnitude of risk and the central role that AI will have in ordering societal infrastructure, those responsible must be taught from the beginning how to design for healthy outcomes. This includes awareness ranging, for example, from data integrity and cultivating transparency to understanding how technical decisions relate to civil, social and geopolitical outcomes.[49]

AI's impact is already seen in our homes, highways, businesses, and professional lives. Today it is embedded in children's toys and classrooms[50]– the time is coming soon when robots will be caring for our children and the elderly. It is essential that policy decisions are made to protect society, particularly the most vulnerable among us. Of course, the challenge is to accomplish this without stifling innovation.

Renowned Australian philosophers, Dr. Matthew Beard and Dr. Simon Longstaff have raised central questions about how, as a world, we approach the use of AI for good and not for bad, noting that this question is not limited to the obvious military or headline-grabbing topics, but must be pursued with respect to those areas where "the stakes aren't obvious and the harms are hard to foresee."[51]

In their joint paper, Beard and Longstaff propose a universal ethical framework for technology based on principles that they argue should inform the design, development and deployment of new technologies regardless of industry sector or AI product. The authors posit that "if ethics frames and guides our collective decision-making" society can enjoy the benefits of AI without falling victim to its avoidable and manageable shortcomings.[52]

Fortunately, AI developers are also recognizing the need to regulate their own activities. Microsoft has developed six principles for its AI systems, including: fairness, safety and reliability, privacy, inclusion, transparency, and accountability.[53] Similarly Google has instituted a "responsible Development of AI" protocol.[54]

But how far can and should self-regulation go with technology that has the inherent ability to violate our rights of privacy and implicate moral dilemmas? Certainly, corporations should not be trusted with the unilateral right to make such decisions for society.

While numerous companies and organizations have been working in good faith on developing codes of ethics to govern the development and deployment of AI, one legal scholar and European Commission official/advisor, Paul Nemitz, sees at least some of these efforts as a "move of genius" by the large tech companies, intended to delay the debate and necessary work on law and regulation over AI. While embracing the concept of codes of ethics when they are intended to guide the behavior of companies above and beyond the rule of law, Nemitz stresses that any effort to replace or avoid the implementation of law through ethics must be rejected. Nemitz observes that the conflicts of interest that inevitably result between corporations and the public cannot be solved by unenforceable codes of ethics or self-regulation.[55] Additionally, there is no guarantee that ethical frameworks developed by current or future AI companies will be compatible with the expectations of those who use the AI or policymakers.

But codes of ethics, backed by enforcement through certifying regulatory bodies, or the violation of which can be proven in a court of law are a step above self-regulation and an approach that should be considered. Just as AI through machine learning is a new frontier of "historical" AI, regulatory approaches need to be innovative and nimble, not simply grounded in regulatory history. A technology-savvy approach to regulation will go far in preventing societal harm; conversely, a stodgy mindset of lawmakers may indeed shut a door we've only just cracked open. Assembly of top minds in AI, untethered to corporate gain, should be at the heart of any regulatory body shaping AI legislation or regulation. And clearly, legislation/regulation should demand full transparency of algorithmic output.

The core question then becomes – which of the challenges that AI presents can be left to ethics and which need to be addressed by standards-setting organizations or by regulation or law. Some, including Nemitz, suggest that those matters that are fundamental rights of individuals or important to the State – must be dealt with through a parliamentary, democratic process. This is not unlike the approach beginning to emerge in the various regulatory arenas.

## VI. TO REGULATE—OR NOT—AND HOW?

Governments and policymakers around the world are grappling with questions about how AI impacts their existing legal and regulatory frameworks. Some obvious and other less obvious threshold questions are being or need to be addressed:

- What interests should regulators seek to protect?
- Do existing regulatory structures cover AI? If not, how should they be adapted, or should new ones be established?
- Recognizing that the cost of regulation is burdensome, how do regulators balance the human benefits to be derived from AI innovation against the potential harm to be protected against?
- How do regulators address these issues in a timely manner recognizing that protections should be in place before consumers and users experience harm, while at the same time knowing that AI is still developing?

- Knowing that entities are investing now in AI, and in some cases, these are regulated entities, is it reasonable to impose regulations after-the-fact that could cause economic harm, or even insolvency?
- Who should regulate AI? Should there be an overarching AI technology regulator, or should the existing sector-specific regulators determine how best to regulate AI within the context of their overall regulatory efforts? In the insurance industry, does this state-by-state approach further exacerbate the effect of extra-territorial regulation that the insurance industry is burdened by now? Or is a central regulator more appropriate to avoid duplication of and inconsistent regulation? Is this decision dictated by the lack of AI skills and human resources in the regulatory arena?

In addition to the issues above, policymakers should carefully consider the question of intentional ambiguity as opposed to specificity in regulating AI. Laws are often enacted with a recognized degree of ambiguity, due to political pressures, stalemate or lack of subject matter knowledge. The expectation is that someone down the road will fill in the details—a regulatory agency may adopt regulations—or a court may interpret and decide what was intended by the legislative body or is required by other laws, treaties or constitutions. Although many commentators have encouraged an "agile" approach, others have suggested that this transfer of responsibility can be counter-productive as applied to AI regulation.[56]

Recognizing that those who build AI systems must by necessity deal with certainty and precision, when faced with ambiguity or broad standards, the tendency will be for data scientists to fill in the gaps. When this occurs, the accountability that we expect through our courts when harm occurs, could be difficult to attain because AI depth is not accessible to the average person and courts are ill-equipped to understand it. These experts suggest that the best approach might be to supplement the more ambiguous or "agile" laws with specific regulation. "Rules give clarity and forewarning while standards offer greater flexibility for interpretation"—in practice, policymakers may consider a broad overarching standard with a narrow set of rules, through regulation, for developers to follow.[57] This is not unlike the federal government's approach to standard-setting in the cybersecurity space.

But once policymakers decide on the degree of desired specificity, should such regulatory standards be sector specific or generally applicable across industries? So far, the sector-specific approach seems to be prevailing.[58] But regardless of the level at which regulation ultimately occurs, it is essential that regulation occur in the context of the application of the AI. The privacy needs that arise from AI systems that recommend songs and movies are very different than the life and death concerns presented by an AI system diagnosing a critical medical condition.

That said, some proponents are contemplating not only a national but a world view of AI. The Boston Global Forum[59] advocates for common standards around the world to promote cooperation and "interoperability" among countries.[60] We need to proceed mindfully and cautiously.

***AI Use by Governments***

While developing their own knowledge and skills in AI and assessing the way in which to regulate their respective industries and protect consumers, governments and regulators are

focused on adopting AI technologies for fraud detection, compliance and other regulatory uses to make their own jobs more efficient and accurate, a topic which is beyond the parameters of this article.[61]

The U.S. House of Representatives has observed that federal, state and local government use of AI to make "consequential decisions" about people should ensure that the algorithms that support these systems are accountable and inspectable.[62] Given the power that government has over our daily lives, not to mention our industries, we need to be ever-vigilant of these developing capabilities and uses.

## VII. AN OVERVIEW OF REGULATORY EFFORTS

### *United States—Executive Order*

On February 11, 2019, President Trump signed an executive order entitled "*Maintaining American Leadership in Artificial Intelligence*."[63] The executive order seeks to solidify American leadership in AI by empowering federal agencies to: drive breakthroughs in AI research and development; establish technological standards to support reliable and trustworthy systems that use AI; provide guidance with respect to regulatory approaches; and to address issues related to the AI workforce.[64] Consistent with the objectives set forth in Section 2 of the executive order, Section 6 directs that within six months the OMB, in coordination with other groups, will issue a memorandum to all agencies that will (1) inform the development of regulatory and non-regulatory approaches by such agencies regarding technologies and industry sectors that are either empowered or enabled by AI, and (2) consider ways to reduce barriers to the use of AI technologies in order to promote innovative application. Both mandates are directed to occur in the context of upholding civil liberties, privacy, and American values and the reduction of AI barriers must be balanced with the protection of the U.S. economy and national security. The agencies have six months form the OMB memorandum to respond to the OMB.

The executive order also directs the National Institute of Standards and Technology (NIST) to issue a plan "for Federal engagement in the development of technical standards and related tools in support of reliable, robust, and trustworthy systems that use AI technologies." The NIST plan is to be developed in conjunction with private sector, academia and non-government entities as appropriate.

NIST has been a leader in the development of cyber security standards. Their work in AI can be expected to be as challenging but as robust an effort as they have led previously.

### *Federal Legislation and Regulation*

To date, there have been six U.S. Congressional hearings held on the topic of AI, the most recent culminated in a report of a subcommittee of the U.S. House of Representatives called *Rise of the Machines*. The House Subcommittee recommended that federal agencies review privacy laws and regulations to determine the extent to which they apply to AI technologies, and where necessary, update existing regulations to accommodate AI. The Subcommittee specifically stressed that any regulatory approach to AI consider whether the risks to public safety or consumers fall within existing regulatory frameworks, and to the extent they don't, consider whether modifications or additions are necessary to better account for AI use.

As part of a multipronged effort, two recent bills have been introduced in Congress. The most interesting one, the Algorithmic Accountability Act, introduced in the Senate with an equivalent bill in the House, would require large tech companies[65] to audit existing machine-learning systems to identify the system impact on accuracy, fairness, bias, discrimination, privacy and security and take corrective action on a timely basis. It would also require such audits to be conducted prior to implementation on new AI systems. The US Federal Trade Commission would be required to promulgate rules within 2 years following enactment of the legislation and would have the power to enforce violations in accordance with existing laws for unfair trade practices and deception. State attorneys general would also have the right to bring civil actions on behalf of the residents of their State with the right of the FTC to intervene.

In Congressional subcommittee hearings, testifying experts frequently cite concerns about individuals' privacy rights, including the risk of breaches by hackers, misuse of personal data by those who collect it, and secondary use concerns—data collected for one purpose is later re-appropriated for another.[66] A difference of opinion exists as to whether Federal privacy regulation is needed for AI or whether regulations should be sector-specific.

Some AI products and practices are currently subject to Federal privacy laws, including HIPAA and the Gramm-Leach-Bliley Act.[67] But Assistant Professor Nicholson Price at the University of Michigan Law School, has pointed out that while HIPAA's Privacy Rule governs and restricts both disclosure and use of certain health information by "covered entities," there are issues with respect to the protection of that information in the AI context, most importantly, that large aggregators of big data, including Google and Apple, are not within the definition of "covered entities." [68]

***U.S. Food and Drug Administration (FDA)***

Professor Price has opined that the typical tools that help ensure safety and efficacy in medical technology—scientific understanding and clinical trials may not work well in the context of AI. Understanding is challenged by the fact that we often don't understand how algorithms make decisions (at least so far), and even if we could the results would likely be too complex. Clinical trials may not be feasible because the algorithms will make highly personalized treatment predictions; the benefits of speed and low cost would be undercut by the long, ponderous, and expensive clinical trial process; and the changing nature of machine learning algorithms would present obvious difficulties.[69]

Professor Price suggests that the FDA focus instead on procedural safeguards such as the quality of the data used, the development techniques, and the validation procedures coupled with robust oversight following their use. The FDA appears to be headed this direction with some of its recent announcements:

- April 11, 2018—the U.S. Food and Drug Administration permitted marketing of the first medical device to use artificial intelligence to detect greater than a mild level of the eye disease diabetic retinopathy in adults who have diabetes.[70]
- May 24, 2018—the FDA approved an algorithm that aids radiologists in detecting wrist fractures.[71]

- October 2018—the FDA outlined five steps it was developing that would affect the regulation of AI in the U.S.[72]
- January 2019—the FDA unveiled its first software pre-certification pilot aimed at streamlining the assessment of safety and effectiveness of software technologies from manufactures that have demonstrated a level of excellence, without inhibiting patient access to those technologies.[73]

***Federal Reserve Board***

The Federal Reserve Board has recognized the need for regulation of AI to be thoughtfully designed so ensure the appropriate mitigation of risks, enhancing efficiency, risk detection and accuracy to the financial sector—but without hindering innovation that would benefit consumers and small businesses.[74] In a presentation Governor Lael Brainard identified several existing regulations, guidance and approaches that will have some applicability to AI regulation, including: (1) the Federal Reserve's "Guidance on Model Risk Management" (SR Letter 11-7) which highlights the importance of embedding critical analysis throughout the development, implementation and use of models, including AI algorithms; (2) examiners' practice of evaluating the processes that firms utilize for developing and reviewing models; (3) existing guidance on vendor risk management (SR 13-19/CA 13-21) and guidance on technology service providers could be expected to apply to externally sourced AI-tools or services; and (4) the regulator's longstanding risk-focused supervisory approach requiring that the level of scrutiny be commensurate with the potential risk posed.[75]

Importantly, the Federal Reserve recognizes that AI presents challenges with respect to opacity and explainability as well as the related issue of the "proverbial black box." Noting that it is not uncommon for questions to arise about the lack of understandability that a bank has regarding its vendors' models, Governor Brainard focused on the need to avoid discrimination and unfair outcomes, ensuring that AI tools do not "learn" the biases of the society in which they were created.[76] He also recognized that the AI community is working on developing "explainable" AI tools.

***The NAIC and State Regulation of Insurance***

With the ability of AI systems to develop independently, without direct human involvement, it will be essential for insurers to develop controls and monitoring that will ensure that machine learning continues to adhere to federal and state regulations concerning data privacy, fairness, discrimination and cybersecurity.

The National Association of Insurance Commissioners (NAIC) is addressing the issue of regulation of AI primarily through two working streams of its Executive Committee—the Innovation and Technology Task Force and the Big Data Working Group. State regulators are grappling with how to obtain and apply the skills necessary to review AI systems going forward with some smaller states preferring that this analysis be done through the NAIC and larger states objecting to the NAIC usurping their role as regulators.

Similar to their federal counterparts, the task force and working group's charges direct them to review existing model laws and regulations in light of AI and recommend changes to accommodate emerging technology. These reviews include marketing, rate regulation,

underwriting, claims, regulatory reporting requirements, the regulation of data vendors and brokers, and consumer disclosure requirements.[77]

The NAIC Casualty Actuarial Task Force adopted a white paper in 2018[78] providing over 90 suggested best practices (guidance) to state regulators for assessing predictive models in the private passenger automobile and homeowners' insurance lines for purposes of rate filing. The task force noted that its guidance was largely transferable to other lines of business and other insurance endeavors (marketing, underwriting and claims). However, it cautioned against using the best practices to create a standard for rate filings in all cases.[79] The American Academy of Actuaries, in reviewing the white paper, suggested that regulators consider whether the guidance has the potential to be unmanageable for modelers as well as regulators.[80]

***New York Regulation***

On January 11, 2018, New York City Mayor Bill de Blasio signed a bill into law that creates a task force to examine how the city's agencies use algorithms to make decisions that can affect millions of New Yorkers. The bill requires NYC establish a task force to recommend ways to establish public accountability for the city's use of algorithms, including ensuring accuracy and fairness. In New York, algorithms are used to assign kids to schools, screen for benefit fraud, assess teacher performance, and design predictive policing behaviors.[81] (It is interesting to note that prior iterations of the bill would have required companies to disclose proprietary algorithms.)

The concern about bias in the insurance context is best illustrated in Insurance Circular Letter No. 1 (2019) issued by the New York Department of Financial Services[82] which followed the Department's review of external data available to insurers for use with "algorithms and predictive models." Citing to New York Insurance laws prohibiting the use of race, color, creed, national origin, status as a victim of domestic violence or past lawful travel, sexual orientation, and other matters, from being used in underwriting,[83] the Department concluded:

> [a]n insurer should not use external data sources, algorithms or predictive models in underwriting or rating unless the insurer has determined that the processes do not collect or utilize prohibited criteria and that the use of the external data sources, algorithms or predictive models are not unfairly discriminatory. The insurer must establish that the external data sources, algorithms or predictive models are based on sound actuarial principles with a valid explanation or rationale for any claimed correlation or causal connection. An insurer must also disclose to consumers the content and source of any external data upon which the insurer has based an adverse underwriting decision.

The Department holds insurers responsible for not just relying on a vendor's claim of non-discrimination or the proprietary nature of their process—the burden remains with the insurer to independently determine the vendor's compliance with anti-discrimination laws.[84] This runs parallel to cyber security law where the insurer cannot avoid third-party failure but is held accountable for that failure.

***European Union***

On April 8, 2019, the European Commission published *Ethics Guidelines for Trustworthy AI* and a document entitled *A Definition of AI: Main Capabilities and Scientific Disciplines*.[85] Based on fundamental rights and ethical principles, the Guidelines list seven key requirements that AI systems should meet and offers guidance for practical implementation of each requirement. A pilot process is established in which stakeholders can participate and provide feedback for improvement, along with a forum to exchange best practices.

The seven key requirements are:

1. Human agency and oversight
2. Technical robustness and safety
3. Privacy and Data governance
4. Transparency
5. Diversity, non-discrimination and fairness
6. Societal and environmental well-being
7. Accountability

The European Commission is considering whether to pursue additional legislation for AI beyond the current General Data Protection Regulation (GDPR), effective May 25, 2018, which applies to AI to the extent it processes personal data. Recognizing that AI regulation must be flexible enough to allow innovation, the Commission is reviewing whether the EU and national frameworks are sufficient for addressing the challenges of AI. Among the factors under consideration are how to make regulation usable and practical, and whether different rules should be applied to public as opposed to private sector uses of AI. A report identifying gaps is expected in mid-2019.

***AI Regulatory Activities in Other Jurisdictions***

On the world stage, the U.S. Congress has recognized the importance of the United States' role in the development and application of AI-driven technologies while at the same time noting that U.S. leadership is no longer guaranteed. Particularly concerning is the prospect of Russia and China overtaking the U.S. in AI advancement and losing any decisive advantage to other nation states. Congress believes that maintaining the U.S.' competitiveness in AI is critical to its economic security.[86]

One of the many organizations working on establishing a code of ethics is the Boston Global Forum which has published the *AIWS Report About AI Ethics*.[87] The report tracks the efforts of the G7 countries (Canada, France, Germany, Italy, Japan, the UK and the US) as well as other influential AI countries (China, India and Russia). The information is not limited to the issue of ethics but reports more generally on AI regulatory and legislative activities.

Additionally, the OECD has several workstreams for AI and has published a number of articles about the impact of AI on labor markets, education/training gaps and more comprehensive scientific uses of AI.[88]

## VIII. HOW SHOULD COURTS AND ARBITRATION PANELS ADDRESS AI?

Although some arbitrators will look at the prospect of deciding cases involving AI as a bridge too far, they might take comfort (or not) in the fact that few in Congress or on the bench understand the foundations of AI any better. As noted above, laws are often ambiguous, and insurance and reinsurance contracts are not always a model of clarity either. Courts and arbitration panels, tasked with the authority and responsibility to determine whether one parties' rights have been violated by the other, or whether a party has departed from a legal standard may be hard-pressed to do so in the context of complicated software systems.

This lack of expertise on the part of courts and panels will shift the determination, or at least, the interpretation, to technical experts. While courts might choose to appoint special masters with special skill and knowledge (determination), arbitration panels and those who appoint them might consider appointment of experts that can help bridge the gap of understanding between the AI system and the panel as decision-makers (interpretation).

The Federal Rules of Evidence set forth the federal standard for the admissibility of new scientific methods in adversarial proceedings: "A witness who is qualified as an expert by knowledge, skill, experience, training, or education may testify in the form of opinion or otherwise if … (a) the expert's scientific, technical or other specialized knowledge will help the trier of fact to understand the evidence or to determine a fact in issue…"

Attorneys presenting cases in arbitration or litigation and decision-makers should take into account that full transparency does not equate with accountability. For example, revealing the underlying mechanics of AI could result in exposing trade secrets, hamper law enforcement or lead to gaming certain decisions. As noted by one expert, concealment does not imply the inability to meaningful analyze the AI system.[89] In a very informative law journal article, Joel Reidenberg sets forth technical tools to provide policy and decision-makers with the ability to protect software trade secrets while at the same time ensuring that AI systems are accountable.[90]

## IX. CONCLUSION

AI will impact our daily lives, our families and our work environments with increasing speed and scope. While we are experiencing only the beginning of these impacts now – governments, policymakers, insurance professionals, courts and arbitration panels will need to internalize the basics of AI as soon as possible – the learning curve can be daunting, requiring us to step outside our comfort zones into the scientific world – but it holds much promise and excitement.

*The author recognizes and thanks Tony Lombardo who provided critical support during research and writing of this article.

---

[1] Matthew Beard & Simon Longstaff, *Ethical by Design: Principles for Good Technology* ("*Ethical by Design*"), The Ethics Centre, at 8. Online at https://ethics.org.au/ethical-by-design/

[2] U.S. House of Representatives, Committee on Oversight and Government Reform, Subcommittee on Information Technology, *Rise of the Machines: Artificial Intelligence and its Growing Impact on U.S. Policy* ("*Rise of the Machines*"), September 2018 at 2. Online at https://www.hsdl.org/?abstract&did=816362

"[AI] has the potential to dramatically reshape the nation's economic growth and welfare." *See also*, "PwC Top Issues," *The Insurance Industry 2015* at 22. Online at https://www.pwc.com/us/en/insurance/publications/assets/pwc-top-issues-the-insurance-industry-2015.pdf. Google estimates that driverless cars may reduce traffic accidents by 90%, reduce the number of cars by 90%, and reduce wasted commute time and energy by 90% -- resulting in a savings of $2 trillion per year to the US economy alone.

[3] Defense Advanced Research Projects Agency (DARPA), "DARPA Announces $2 Billion Campaign to Develop Next Wave of AI Technologies," (Sept. 7, 2018). Online at https://www.darpa.mil/news-events/2018-09-07.

[4] "Artificial Intelligence: Emerging Opportunities, Challenges, and Implications*" GAO Technology Assessment*, (March 2018). Online at https://www.gao.gov/assets/700/690910.pdf.

[5] Joel Reidenberg, "Accountable Algorithms," *Fordham Law School* (2017) at 633. Online at https://ir.lawnet.fordham.edu/cgi/viewcontent.cgi?article=1888&context=faculty_scholarship. *See also*: *Rise of the Machines* at 4; and "Artificial Intelligence Collides with Patent Law," *World Economic Forum* (2018). Online at http://www3.weforum.org/docs/WEF_48540_WP_End_of_Innovation_Protecting_Patent_Law.pdf. The Creative Machine developed in 1994 by AI pioneer Stephen Thaler, is known for having created the first known patented AI-generated invention. Computers have also designed a new nose cone for a Japanese bullet train, novel piston geometries for reducing fuel consumption in diesel engines and new pharmaceutical compounds.

[6] *What are we Learning about Artificial Intelligence in Financial Services?*" ("AI in Financial Services") (November 13, 2018) Presentation by Federal Reserve Board Governor Lael Brainard, Fintech and the New Financial Landscape Conference, Philadelphia, PA.

[7] This type of machine learning is called "supervised learning:" the algorithm is fed a set of training data that contains labels on some items in the training set. The algorithm will "learn" a general rule of classification that it will use to predict the labels for the remaining items in the set. For example, after being provided with labels that show various features of dogs, the AI will be able to distinguish a dog from a cat, or a turtle or a human.

[8] This type of machine learning is referred to as "unsupervised learning:" the data provided to the algorithm does not contain labels. The algorithm is asked to detect patterns in the data by identifying clusters of observations that depend on similar underlying characteristics. For example, an AI looking at dogs might focus on those beings that have tails. When it erroneously labels a cat to be a dog because the cat too, had a tail, the developer might feed images of cats and dogs back through the algorithm to find other distinguishing patterns and thereby "learn" the difference between cats and dogs. This feedback is referred to as "reinforcement learning."

[9] David Silver, et.al. "AlphaGo Zero: Learning from Scratch," *DeepMind Technology*. Online at https://deepmind.com/blog/alphago-zero-learning-scratch/

[10] Aaron Frank, "Machines Teaching Each Other Could be the Biggest Exponential Trend in AI," *SingularityHub* (January 21, 2018). Online at https://singularityhub.com/2018/01/21/machines-teaching-each-other-could-be-the-biggest-exponential-trend-in-ai/#sm.0000jdet2232zd0tx1g2cctlu7h0f. "There will be a big leap in unsupervised learning in the near future. We will see companies using AI to train I. Instead of data scientists...companies will let AI do the work for them." Yaron Hadad, Co-founder & Chief Scientist, Nutrino.

[11] "In this new line, new players that have generated deep risk insights, are also expected to enter these unpenetrated segments of the market: for example, life insurance for individuals with specific diseases." "PwC Opportunities Await: How InsurTech is Reshaping Insurance," ("PwC Opportunities Await") *Global Fintech Survey* (June 2016) at 10. https://www.pwc.lu/en/fintech/docs/pwc-insurtech.pdf.

[12] PwC Opportunities Await at 7. "Clients now expect personalized insurance solutions, and "one-size" does not fit all." Product innovation might include, for example: usage-based, driving mode-based, and trip-based insurance using telematic devices and automated driver assistance systems (ADAS). Opportunities might also include unbundling current auto insurance and re-bundling them with products targeted to urban, casual and self-driving insured profiles. PwC Top Issues at 22.

[13] PwC Opportunities Await at 7. Linking crowdsourcing with insurance is the P2P insurance business concept of the sharing economy. Crowdsource insurance provides access to effective risk capital by bringing together a pool of policyholders.

[14] "The rise of affinity groups, car-sharing groups, and vehicle manufacturers who want to package auto insurance with autonomous vehicles can open up new distribution channels for auto insurers." PwC Top Issues at 22.

[15] PwC Opportunities Await at 10. "Current trends show an increasing interest in finding new underwriting approaches based on the generation of deep risk insights…Initially incumbents viewed UBI [usage-based insurance] as an opportunity to underwrite risk in a more granular way by using new driving/behavioural variables, but new players see UBI as an opportunity to meet new customer needs (e.g. low mileage or sporadic drivers)." Risk-slicing (see Note 16) introduces new pricing opportunities for insurers such as usage-based or mileage-based premium. Similarly, alternating between hands-on and hands-off driving (see Note 16) may suggest a different approach to pricing based on the mode of driving. PwC Top Issues at 21.

[16] PwC Opportunities Await at 10. Car-sharing continues to grow. Particularly popular with urban millennials, car-sharing memberships are expected to reach 26 million by 2020. In the next 5-10 years we are likely to see an increase in cars incorporating a self-driving mode. Depending on road conditions and personal preferences, a portion of a single trip could be human hands-on driving and other portions could be hands-off driving. Thus, a single trip could trigger different liabilities – driver liability for the hands-on portion and product liability for the hands-off portion. PwC Top Issues at 21. Both the car-memberships and the different modes of driving are examples of "risk slicing."

We will soon see more cars with a self-assisted/driving mode allowing a driver to shift between hands-on and hands-off driving. This could result in different risk profiles for a single trip and premiums that are priced differently depending on the mode of driving.

[17] "Adversarial machine learning" involves feeding information into a machine learning system causing it to misbehave, devising ways to deceive it or to reveal sensitive information. *See* "How Malevolent Machine Learning Could Derail AI," *MIT Technology Review* (March 25, 2019). Online at https://www.technologyreview.com/s/613170/emtech-digital-dawn-song-adversarial-machine-learning/

[18] The implications of autonomous car technologies will significantly impact the auto insurance sector, including lower premiums. PwC Top Issues at 18.

[19] "Machine Learning in Insurance," *Accenture FS Perspectives* (2018). Online at https://www.accenture.com/_acnmedia/PDF-84/Accenture-Machine-Leaning-Insurance.pdf

[20] "Intelligent Machines and the Transformation of Insurance," *NAIC & The Center for Insurance Policy and Research CIPR Newsletter* (January 2019) at 15. Online at https://www.naic.org/cipr_newsletter_archive/vol26_intelligent_machines.pdf

[21] Curt Hopkins, "4 Obstacles to Ethical AI (And How to Address Them)" ("HP 4 Obstacles"), *Hewlett Packard Enterprise*. Online at https://www.hpe.com/us/en/insights/articles/4-obstacles-to-ethical-ai-and-how-to-address-them-1807.html. "AIs are created and trained by people. If the pool of developers and trainers are similar in background, their shared bias is likely to be communicated to—and then used by—the AI . . . The more diverse the group, the less likely an AI system is to reinforce shared bias."

[22] "This is How AI Bias Really Happens—And Why It's So Hard to Fix," ("How AI Bias Really Happens") *MIT Technology Review* (February 2019). Online at https://www.technologyreview.com/s/612876/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/?utm_campaign=the_algorithm.unpaid.engagement&utm_source=hs_email&utm_medium=email&utm_content=71709273&_hsenc=p2ANqtz-86eu1cH64EtSPH7aKVvYPFWoIDw-oi1cBaeJ6xQsAprf7_vGNkIMfmSa3nAwfvHRVTYtk08dHC6knRRAlT2lRCMgCMbA&_hsmi=71709273

[23] How AI Bias Really Happens.

[24] *Id.*

[25] For example, certain zip codes may suggest lower income and be a proxy for racial bias.

[26] HP 4 Obstacles, "Those issues go beyond gender and race; they also encompass what you studied, the economic group you come from, your religious background, all of your experiences."

[27] HP 4 Obstacles.

[28] How AI Bias Really Happens.

[29] *Id.*

[30] *Rise of the Machines* at 11.

[31] *Id.* at 10-11.

[32] "Deep Learning" has been defined by the Financial Stability Board as "a form of machine learning that uses algorithms that work in 'layers' inspired by the structure and function of the brain. Deep learning algorithms, whose structure are called artificial neural networks, can be used for supervised, unsupervised, or reinforcement learning." *Artificial Intelligence and Machine Learning in Financial Services*, ("*AI and ML in Financial Services*") Financial Stability Board (November 1, 2017) at 5. Online at http://www.fsb.org/wp-content/uploads/P011117.pdf

[33] "The Dark Secret at the Heart of AI," *MIT Technology Review* (April 11, 2017). Online at https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/

34 *Science, Technology and Innovation Outlook 2018,* OECD, at 132, Box 5.4. Online at https://read.oecd-ilibrary.org/science-and-technology/oecd-science-technology-and-innovation-outlook-2018_sti_in_outlook-2018-en#page134

[35] *AI and ML in Financial Services* at 33-34.

[36] *Id.* at 34.

[37] *Id.*

[38] Paul Voosen, "The AI Detectives," *Science* (July 7, 2017). Online at https://science.sciencemag.org/content/357/6346/22/tab-pdf

[39] *Rise of the Machines* at 12.

[40] Sean Gallagher, "Researchers Scared by their Own Work, Hold Back 'Deepfakes for Text' AI," ARS Technica (February 15, 2019). Online at https://arstechnica.com/information-technology/2019/02/researchers-scared-by-their-own-work-hold-back-deepfakes-for-text-ai/. "OpenAI, a non-profit research company investigating "the path to safe artificial intelligence," has developed a machine learning system called Generative Pre-trained Transformer-2 (GPT-2), capable of generating text based on brief writing prompts. The result comes so close to mimicking human writing that it could potentially be used for "deepfake" content…The performance of the system was so disconcerting, now the researchers are only releasing a reduced version of GPT-2…"

[41] Malicious Deep Fake Prohibition Act of 2018, S.3805, 115th Cong. 2018, Senator Ben Sasse (R-NE). Online at https://www.congress.gov/bill/115th-congress/senate-bill/3805/text

[42] NY Assembly Bill A08155 and companion bill S05857, introduced May 31, 2017, online at https://nyassembly.gov/leg/?default_fld=&leg_video=&bn=A08155&term=2017&Summary=Y&Text=Y. The summary of the bills provide: "[e]stablishes the right of privacy and the right of publicity for both living and deceased individuals; provides that an individual's persona is the personal property of the individual and is freely transferable and descendible; provides for the registration with the department of state of such rights of a deceased individual; and that the use of a digital replica for purposes of trade within an expressive work shall be a violation." *See also,* https://www.theregister.co.uk/2018/06/12/new_york_state_is_trying_to_ban_deepfakes_and_hollywood_isnt_happy/

[43] "Deep Fakes: Let's Not Go Off the Deep End," *TechDirt*, (January 30, 2019). Online at https://www.techdirt.com/articles/20190128/13215341478/deep-fakes-lets-not-go-off-deep-end.shtml.

[44] "The Adversarial Robustness Toolbox: Securing AI Against Adversarial Threats," *IBM Research Blog* (April 17, 2018). Online at https://www.ibm.com/blogs/research/2018/04/ai-adversarial-robustness-toolbox/

[45] "AI and Adversarial Attacks: Vulnerabilities to Manipulation," *Harvard Magazine* (January-February 2019). Online at https://harvardmagazine.com/2019/01/ai-and-adversarial-attacks

[46] HP 4 Obstacles.

[47] *Empowering AI Leadership*, World Economic Forum. Online at https://www.weforum.org/projects/ai-board-leadership-toolkit

[48] *Teaching AI Ethics*, World Economic Forum. Online at https://www.weforum.org/projects/teaching-ai-ethics

[49] *Id.*

[50] *Generation AI: What Happens when your Child's friend is an AI Toy that Talks Back?* World Economic Forum,. Online at https://www.weforum.org/agenda/2018/05/generation-ai-what-happens-when-your-childs-invisible-friend-is-an-ai-toy-that-talks-back/

[51] *Ethical by Design* at 8.

[52] *Id.*

[53] "Why the Government Must Help Shape the Future of AI," *Brookings Institute* (September 13, 2018) at footnote 1. Online at https://www.brookings.edu/research/why-the-government-must-help-shape-the-future-of-ai/

[54] *Id.* at footnote 2.

[55] Nemitz, Paul Friedrich, *Constitutional Democracy and Technology in the Age of Artificial Intelligence* (August 18, 2018). DOI 10.1098/RSTA.2018.0089 - Royal Society Philosophical Transactions A. Available at SSRN: https://ssrn.com/abstract=3234336 or http://dx.doi.org/10.2139/ssrn.3234336.

[56] Accountable Algorithms at 700.

[57] *Id.* at 702.

[58] "Will the UK Regulate AI?" *Slaughter and May* (July 27, 2018). Online at https://www.slaughterandmay.com/media/2537013/will-the-uk-regulate-ai.pdf, noting that the House of Lords has suggested an existing sector-specific approach to AI regulation. While the UK government appears to agree with this

approach, it is establishing a Ministerial Working Group on Future Regulation three new AI organizations, including: The Centre for Data Ethics and Innovation (CDEI), The AI Council and The Government Office of AI.

[59] The Boston Global Forum includes the Michael Dukakis Institute for Leadership and Innovation which sponsors the AI World Society ("Boston Global Gorum"). Online at https://dukakis.bostonglobalforum.org/.

[60] Boston Global Forum.

[61] AI in Financial Services at footnote 8. AI tools are likely to be useful for central banks and regulators relating to supervision, financial stability and monetary policy; noting that AI could be useful to central banks and prudential regulators for applications ranging from systemic risk identification to detecting fraud and money laundering. Financial Stability Board, Artificial Intelligence and Machine Learning, 2017. The State of Ohio uses robotics in its criminal investigation bureau to help reduce the turnaround time on untested rape kits, facilitating the testing of 14,000 untested rape kits and identifying 300 serial rapists linked to 1,100 crimes. *Rise of the Machines* at footnotes 13 and 14.

[62] *Id.* at 1, 10.

[63] Executive Office of the President, Maintaining American Leadership in Artificial Intelligence, Executive Order 13859. Online at https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence

[64] Jadzia Pierce & BJ Altwater, "President Trump Signs Executive Order on Artificial Intelligence," *Cov Financial Services* (February 12, 2019). Online at https://www.covfinancialservices.com/2019/02/president-trump-signs-executive-order-on-artificial-intelligence/

[65] Algorithmic Accountability Act, Section 2 Definitions, (5) Covered Entity includes greater than $50 million gross receipts and over 1 million consumers or consumer devices and other factors.

[66] *Rise of the Machines* at 9. These fears were realized when one of the nation's largest credit reporting agencies, Equifax, Inc. was breached in 2017 and hackers obtained personal data on 14.5 million Americans.

[67] *Id.*at 10. The Federal Trade Commission is the primary Federal privacy regulatory body and has jurisdiction over other privacy laws and regulations.

[68] "Artificial Intelligence in Health Care: Applications and Legal Implications" ("AI in Healthcare"), W. Nicholson Price, II, *University of Michigan Law School* (2017) at 10-11. Online at https://repository.law.umich.edu/cgi/viewcontent.cgi?article=2932&context=articles. Professor Price teaches intellectual property, health law and innovation in life sciences. Professor Price has a particular interest in regulation of emerging technology. https://www.nicholsonprice.org/

[69] AI in Healthcare at 12-13.

[70] "FDA Permits Marketing of Artificial Intelligence-based Device to Detect Certain Diabetes-Related eye Problems," *FDA News Release* (April 11, 2018,). Online at https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm604357.htm

[71] Food and Drug Administration, "FDA Permits Marketing of Artificial Intelligence Algorithm for Aiding Providers in Detecting Wrist Fractures" (May 24, 2018). Online at https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm608833.htm."The OsteoDetect software is a computer-aided detection and diagnostic software that uses an artificial intelligence algorithm to analyze two-dimensional X-ray images for signs of distal radius fracture, a common type of wrist fracture… It is an adjunct tool and is not intended to replace a clinician's review of the radiograph or his or her clinical judgment."

[72] "5 Ways the FDA Promises to Upgrade AI-Related Medical Devices," *Radiology Business* (November 6, 2018). Online at https://www.radiologybusiness.com/topics/policy/how-fda-will-regulate-ai-related-medical-devices

[73] "FDA Unveils AI Pre-Certification Program," *AI Powered Healthcare* (January 15, 2019). Online at https://www.healthcareitnews.com/ai-powered-healthcare/fda-unveils-ai-pre-certification-program-0. See also, U.S. Food & Drug Administration, "Developing a Software Precertification Program: A Working Model" (January 2019). Online at https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/DigitalHealthPreCertProgram/UCM629276.pdf

[74] AI in Financial Services.

[75] AI in Financial Services.

[76] AI in Financial Services.

[77] NAIC, Innovation and Technology (EX) Task Force 2019 Charges, online at https://www.naic.org/cmte_ex_ittf.htm

[78] *Regulatory Review of Predictive Models*, NAIC Casualty Actuarial task Force, October 25, 2018, online at https://www.naic.org/documents/cmte_c_catf_exposure_predictive_model_white_paper.pdf

[79] *Id.* at 4.

[80] Letter from Richard Gibson to Kris DeFrain, *American Academy of Actuaries* (January 22, 2019). Online at https://www.actuary.org/sites/default/files/files/publications/CASTF_Predictive_Modeling_Comments_20190122.pdf

[81] "New York City's Push for Accountable Algorithms," *Fordham Intellectual Property Media & Entertainment Law Journal* (April 14, 2018). Online at http://www.fordhamiplj.org/2018/04/04/new-york-citys-push-for-accountable-algorithms/

[82] New York Department of Financial Services, "Insurance Circular Letter No. 1" (January 18, 2019). Online at https://www.dfs.ny.gov/industry_guidance/circular_letters/cl2019_01

[83] *Id.* noting Insurance Law Article 26 and Laws §§ 4224(a) and (b)(2).

[84] *Id.* The Department referenced the "disparate impact" standard, seemingly applying it to life insurance, but that topic is beyond the scope of this paper.

[85] Both documents were prepared by The Commission's High-Level Expert Group on AI. The independent group consisting of 52 experts from academia, business and civil society, was appointed in June 2018 following the publication of the Communication from the Commission to the European Parliament, the European Council, the Council, The European Economic and Social Committee and the Committee of the Regions on Artificial Intelligence for Europe on April 25, 2018 (sometimes referred to as the "European Approach on AI").

[86] *Rise of the Machines* at 4-6.

[87] Boston Global Forum.

[88] "Artificial Intelligence," OECD. Online at http://www.oecd.org/going-digital/ai/. *See also*, http://www.oecd.org/going-digital/ai/oecd-initiatives/

[89] Accountable Algorithms at 705.

[90] *Id.*