2022 CLM Focus

November 2-3rd

Washington, DC

"Artificial Intelligence and the IoT Gone Rogue: A Cautionary Tale"

I. Introduction:

As technology continues to rapidly advance, the use of artificial intelligence ("AI") and internet of things ("IoT") devices have become a more integral part of our daily lives. However, with the exponential use of these technologies comes an increase in the liability risks to consumers, programmers, and manufacturers alike. Numerous poorly developed AI have shown patterns of discrimination and proposed flawed or even life-threatening solutions to problems. IoT devices, have continued to injure and kill people when they malfunction while also being a major cybersecurity risk. With the rapid adoption of these technologies, the legal and insurance industries have struggled to deal with the new and unanticipated risks that accompany them. This article addresses the emerging risks associated with unchecked and unexpected consequences of artificial intelligence and IoT going rogue and how claims professionals can best prepare to handle these claims.

II. Emerging Sources of Liability

a) Artificial Intelligence

Both businesses and individuals have become increasingly reliant upon AI and current trends show that we are only going to become more dependent as time goes on. In the McKinsey Global survey conducted in 2021, 56% of businesses reported the adoption of AI to perform operation functions. Businesses are not the only ones employing AI, individuals and governments have both realized the potential AI presents. In 2021, private investment in AI totaled $93.5 billion and more than doubled the private investments made by individuals in 2020. Additionally, federal lawmakers in America have started to notice the impact AI can have. While, in 2015, there was only one bill that proposed regulations involving AI, 130 were proposed in 2021. This increase of attention shows that more people are becoming aware of the prospects AI has to offer as well as the threats that AI can pose if not properly controlled. However, as the prevalence of AI continues to increase in our society, so too will the risk of liability caused by AI malfunctioning.

Microsoft's "Tay" chatbot offers a cautionary tale of an AI system gone rogue. Tay "learned" from Twitter users to make extremist and bigoted statements. Amazon scrapped a

machine-learning tool for selecting the top candidate resumes because it discriminated against women. The most recent artificial intelligence gaffe was Google's facial recognition app labelled a black couple as "gorillas". While these incidents caused embarrassment and outrage over the conclusions these faulty AI produced, these damages are still of the lower end of the spectrum of the harm AI can produce.

Alarmingly, recent research has shown that AI bias may also work its way into robotics. The experiment was conducted by giving a virtual robot instructions to place one of two blocks with a person's face on it into a box, if the block correlated with the command. For example, the robot may have been commanded to "place the criminal block in the brown box" and the robot would decide which blocks were "criminals" to sort into the box. The study showed that, compared to white male blocks, the robot labeled black men as "criminals" 9 percent more and Latino men as "janitors" 6 percent more. Additionally, black and Latina women were more often labeled as "homemakers" and all women were less often categorized as "doctors" by the robot than men. The researchers pointed out that for many of the commands the correct answer would have been for the robot to do nothing because there was not enough information for the robot to categorize the blocks. If, for example, the robot was told to decide which block was a doctor, it should not have done anything because there would not be enough information to conclude either was a doctor. However, for most of those commands, the robot decided to categorize the block based solely on the person's appearance. The robot only correctly categorized no block one third of the time. This study gives cause for concern for the future of robotics as the automation industry is expected to grow to $60 billion by the end of the decade and robotic labor increases. There is also a fear that robots will cause more severe damages based on their bias as they become capable of increasingly complex tasks. As such, it is very possible that companies will have a duty to monitor the actions of their AI driven robots to prevent these biases and, accordingly, be found liable for damages caused by a robot's bias if not fixed.

Other instances of faulty AI have had the potential to ruin individuals' lives or cause serious bodily harm or death. For example, many law enforcement agencies now use AI to help them find criminals. However, such AI often misidentifies people as criminals and can lead to false arrests and put them in danger. Amazon's Rekognition AI was one of these faulty AI that misidentifies twenty-seven professional athletes as criminals during a test run by the ACLU that compared pictures of people with criminal mugshots. Fortunately, this was only a test, so the AI was not actually being used to arrest people. However there have been cases where AI was used to arrest innocent people. In June 2020, police in Detroit relied on a facial recognition AI to make the arrest of Robert Julian-Borchak Williams for felony larceny. The program compared a blurry picture from surveillance video of the real perpetrator shoplifting $3,800 worth of timepieces. As a result, Williams was arrested in front of his wife and children and wrongfully detained. Although Williams was eventually able to prove his innocence, this incident shows the potential for people to be endangered by the mistakes of AI.

AI bias may also be contributing to systemic racism in the criminal justice system because AI algorithms are now being used in a number of states to assist judges in making sentencing and bail decisions. In a 2016 study by ProPublica, one such AI was twice as likely to incorrectly label black prisoners as being at high risk of reoffending as white prisoners. While the race of the prisoner was not directly considered by the AI, the other variables that were considered clearly disfavored blacks. As a result, many black prisoners may be getting stricter jail sentences or higher bail because these AI are incorrectly labeling them.

Racial bias of AI has also affected the administration of medical care. In October 2019, researcher at UC Berkley discovered that black people were receiving lower health risk scores than white people despite being at higher risk of health problems in reality. This was a problem because the risk score the AI gave a patient would affect what standard of care the patient would receive. As a result, many black people received medical care that was below the standard they should have received for someone in their condition. The researchers discovered that the AI incorrectly assumed that the more money that was spent on an individual's healthcare meant that that person was more likely to be at risk of health issues. Eventually, the researchers were able to correct the bias by making the AI focus on different variables, such as avoidable costs and the number of chronic conditions a patient needed treatment for in a year. However, as AI continue to make resource allocation decisions going forward, these biases can lead to severe liability issues when they deny protected classifications of people lifesaving goods and services.

Errors made by AI also could have direct lethal consequences, as shown by IBM's AI, Watson. In 2017, Watson was used to recommend treatments for cancer patients. In one case involving a 65-year-old lung cancer patient that developed severe bleeding, Watson recommended giving the patient a drug that could cause a fatal hemorrhage. Fortunately, the doctor supervising the AI knew better and refused to give the patient the medication. However, as AI becomes more prevalent in our society there will surely be cases where professionals will incorrectly default to an AI's recommendations and severely injure others.

For these reasons, it is clear that AI left to run unchecked can represent major and even existential risks to insureds. Accordingly, claims professionals will need to prepare for the liability threats that AI reliance may cause their clients.

b) Internet of Things

IoT is used to describe devices that have sensors, software, and other technologies that connect and exchange data with other devices over the internet. While this connectivity has led to many improvements to our quality of life, there are also inherent risks that the IoT presents. By 2025, 75 billion active IoT devices will be connecting to the Internet. Accordingly, claims professionals can expect to see claims ranging from property damage and business interruption due to threat actors taking down the grid to wrongful death and catastrophic injury claims.

i) Dangers of Autonomous Vehicles

When IoT devices malfunction, they have the potential to wreak havoc and create unexpected liability exposures. Cautionary tales include everything from rogue crop dusting IoT devices destroying crops to semi-automated vehicle failures resulting in serious accidents. In the last couple of years alone, people have been involved in numerous crashes because of their reliance on the self-driving capabilities of newer vehicles. The Tesla autopilot feature, for example, has become the poster child of this issue. Ever since Tesla sold cars with their autopilot feature, people have been getting into crashes as they metaphorically and literally fell asleep at the wheel. Fortunately, most crashes have been relatively minor up to this point, but there have been more severe accidents that resulted in serious injury or even death. For example, a man in California is currently facing two counts of vehicular manslaughter for crashing a Tesla Model S that was on autopilot into another car. Police reported that the vehicle left a freeway at highspeed before it ran the red light and collided with the victims. With an estimated 765,000 Tesla vehicles in the United States that have similar autopilot functions and even more vehicles with a similar feature, claims professionals should be prepared to see an increase in crashes involving such malfunctions.

Tesla is not the only company that is trying to profit from self-driving cars with deadly results; Uber is in the process of developing autonomous cars for their ride-sharing app. Similarly, Uber's self-driving car has not been without controversy. In fact, one of Uber's test drivers for their fully autonomous vehicles is now facing negligent homicide charges for allowing the car to hit and kill a pedestrian who was walking her bike across the road. Investigations revealed that the test driver who was supposed to be watching the road was streaming a television show instead. Although Uber was not held criminally liable in this situation, they received heavy criticism that caused them to halt their autonomous vehicle testing. However, this event did not deter them for good. In 2020, Uber's self-driving cars were allowed to continue testing in California, along with 65 other transport firms, and have shown a commitment to improving and implementing their autonomous cars. As such, claims professionals should prepare to see more accidents caused by distracted drivers.

ii) Robotic Revenge

Unfortunately, self-driving cars are not the only examples of IoT that endanger people when they malfunction. Several devices in the medical field have the potential to severely injure or kill patients if they go haywire. Malfunctioning robotic surgery systems, for example, have caused several deaths and injuries. In 2016, a study found 144 deaths, 1,391 patient injuries, and 8061 device malfunctions related to robotic surgical machines from 2000 to 2013. The reported list of malfunctions during surgeries included uncontrolled movements of the robot, spontaneously switching on and off, loss of video feed, system error codes, and electrical sparks causing burns. Overall, the study found 550 adverse surgical events per 100,000 surgeries.

Robotics in manufacturing plants have also proven to be deadly when they malfunction. There have been many instances where workers have been crushed and bludgeoned by robotics in these facilities when these robots go haywire. Often these accidents are caused by the robots' strict adherence to their coding, so when a robot encounters a novel situation, it often does not know how to respond. Such a case occurring in Michigan on March 15, 2017, when a maintenance technician named Wanda Holbrook got her head pinned by a robot attempting to move a part down an assembly line while another robot attempted to weld it. As a result, Holbrook experienced severely burns to her face, nose, and mouth and was pronounced dead by first responders at the scene. In this case, Holbrook was performing maintenance on a separate section of robots adjacent to where the incident occurred. When she requested entry into that section, the entire assembly line stopped as the robots in other sections would not continue their work if their sensors told them the next section has not first finished its task. However, only the robots in her section were actually deactivated; the rest were waiting for their sensors to tell them to continue working. When she entered, Holbrook noticed that one of the robots she was to work on was reaching into the previous section to move a part to its section, so she entered the previous section to manually guide the robot to its section. Unfortunately, doing this convinced the previous sections that the part they were working had progressed to the next section, causing the robots in the incident section to continue work and subsequently kill Holbrook. Automation of many factory jobs is an inevitability at this point, and with the increase in automation, there will also be an increase in maintenance workers. Therefore, events like this will become more common going forward and claims professionals will have to know how to deal with them.

There is no denying that AI and IoT devices will become more and more integrated into our daily lives as time goes on. However, as that happens, new and unanticipated risks will begin to emerge alongside them. Therefore, it has become abundantly clear that claims professionals will have to adapt to the changing times to understand how to best handle the damages that result from the technological advancement of society.

III. Legal Response to AI and IoT

As AI and IoT are relatively new fields of technology that have only been widely commercially available for the last few years, there is not much established law regarding liabilities caused by them. Accordingly, there is no established consensus regarding how damages caused by these technologies should be handled. However, that has not stopped courts and legal scholars from developing their own legal theories for the allocation of liability stemming from AI and IoT.

c) Legal Theories for AI Liability

AI presents particularly unique legal liability challenges because, in theory, the software program that was sold to the user will not be the same program that caused the liability. This is because the machine learning capabilities of AI necessarily results in the software rewriting its own code to evolve with the data that it is receiving and become more efficient or accurate in

performing its task. As such, the legal community is debating who should be held liable for damages caused by AI. Is the consumer that used the AI liable because the defect manifested itself while under his control or should the manufacturers and programmers have predicted the defect and prevented it from manifesting preemptively? If the suppliers are to blame, which link in the chain of production is at fault for which percentage of damages? Can an AI be treated as a legal entity — like corporations are — and be held directly responsible for damages it causes? To answer these questions, several legal theories have been proposed to allocate liability for damages caused by AI.

One legal theory proposed is to use the doctrine of Respondeat Superior for AI liability. Respondeat Superior, also known as the "Master-Servant Rule," states that a principal should be liable for the actions of an agent who was negligent while working within his scope of employment. The doctrine was developed in ancient Rome as a way impose liability upon a slave owner for a tort committed by a slave. Despite the unfortunate circumstances that led to the development of this doctrine, some scholars are drawing parallels between AI and Roman slaves to reason why Respondeat Superior should be used to allocate AI liability. Like the Roman slaves, AI are property rather than subjects of law and are capable of making their own decisions that may lead to damages. Therefore, scholars that follow the Respondeat Superior theory believe that the owners of the AI should be liable for torts committed by the AI. This may be either the consumer who bought the AI or the developer who licensed use of the AI.

A similar theory is to treat AI as if it was an ordinary tool. Just as a person would be liable for negligently operating heavy machinery, an AI consumer would be liable for negligently implementing AI. Under this theory, a person who buys or uses the AI will be liable for damages caused by the AI while acting under the person's control. This creates a principal-agent like relationship between the consumer and the AI so that actions taken by the AI could be imparted upon the consumer. Therefore, the consumer would not be able to evade liability for damages caused by the AI simply because he did not intend for the AI to act in the way it did. Under this theory, the AI user could hold the AI developer liable for damages only if he can prove the AI was defective while under the developer's control and that defect was the proximate cause of the damages.

Alternatively, some believe that AI should be treated as a legal entity, similar to corporations. The argument for this position is that AI are capable of rational thought and independent decisions, so they should be held liable for damages they cause. The benefit of this approach is that it becomes easy to identify the liable party because the AI itself will be directly suable when it malfunctions instead of the user or creators who could not have anticipate the failure. However, opponents of this theory argue an AI's rational processes are not equivalent to full legal capacity and AI lack free thought since they must still act within the parameters of their code. Therefore, the opposition concludes that AI lack the mental capacity to be a legal entity. Another issue with this theory is that it limits the recourse available to claimants damaged by AI

malfunctions. Most notably, it is unlikely that AI alone will have enough resources to adequately compensate any claimants it has damaged. Lastly, it is also worth mentioning that legislators are apprehensive to grant AI legal personhood status. For example, the European Parliament has already rejected granting AI legal entity status three times because AI lack legal personalities and human consciences. Accordingly, it is unlikely that AI will be given legal entity status any time soon.

Perhaps the most likely approach that the courts will take for AI liability is to adapt the current laws for product liability to AI. This is because product liability law has historically had to evolve alongside emerging technologies in the past, so it is likely to evolve with the emergence of AI to address novel issues. If that is the case, an AI user could be found liable for the damages caused by an AI if he used the AI in a negligent manner to cause the damages. However, a user would not be liable for damages if he used the AI in a reasonably foreseeable way that inadvertently causes the AI to develop a defect. In such a case, software companies that developed the AI would likely be found liable under product liability law for failing to anticipate how an AI could develop defects through reasonably foreseeable interactions with humans. Courts would likely perform a risk-utility test to determine if the safety precautions could have been taken to decrease risk of AI malfunctions without lowering the AI's utility or unnecessarily increasing the cost of producing the AI.

d) Liability for IoT Damages

IoT claims landscape is equally complex. As the cybersecurity risks of IoT becomes more prominent with the rise of IoT, courts will also likely be more willing to find companies that produce IoT liable for products that lack adequate security measures. Product liability arises when a product was in a defective condition while under a producer's control that made it unreasonably dangerous and was the proximate cause of a plaintiff's damages. People are now becoming aware of just how vulnerable many IoT are to being hacked so, courts may determine that the lack of security on these devices is a defect that was present while under the producer's control and the cause of a user's damages if hacked into. As a result, many producers would then be liable for damages that result from a breach of data caused by their IoT products. However, IoT producers may not be the only ones held liable for IoT cybersecurity issues; IoT users will likely also face liability.

Most companies today use IoT in some capacity for their day-to-day operation. While conducting business, they will inevitably collect sensitive customer data that they have a duty to protect. The problem is many IoT devices that are vulnerable to hacking are not monitored and their software is never updated. Unsurprisingly, these devices often get hacked and act as backdoors for hackers to gain more access to sensitive customer information. As these cyberattacks become more frequent, courts will likely start holding companies at a higher standard of care to take proper precautions in ensuring all IoT devices connected to a network are secure.

Therefore, negligence claims against companies that have substandard IoT cybersecurity will likely increase in the years to come.

IV. Claims Professional Response to AI and IoT Liability

With an increase in AI products, claims professionals will need to make a fundamental shift in the processing and evaluation of claims. These claims will require far more technological sophistication. The claims handler will be well served by developing a deep understanding of technology and approaching A/I and IoT claims like complex have to prepare product liability claims as opposed to simple negligence cases. This is because any accident that involved the product could have been caused by its AI. Claims professionals will have to be prepared to follow the chain of production for any AI sold to determine which point of the manufacturing process may have been responsible for the damages. Therefore, it will be crucial for claims professionals to find experts for various types of AI to analyze claims and determine if the AI malfunctioned and who is to blame if it did. Additionally, claims professionals that cover producers of AI products will need to adjust their rates based on how predictable the AI's behavior is and the products potential to cause damages if the AI malfunctions.

The evolution of technology necessarily results in the evolution of insurance products. New insurance products are already being developed to respond to the risks associated with artificial intelligence and IoT devices. Claims professionals will need to keep abreast of the insurance product iterations to conduct a proper coverage analysis at the outset of the claims handling process.

Like with AI products, claims professionals will also need to gather new resources and experts to evaluate the unique dangers IoT devices present. Claims professionals will not only need to be able to tell if an IoT device's programming was the cause of damages in a claim, but also if a lack of cybersecurity caused the damages. Furthermore, because any company could be liable for a cybersecurity breach, claims professionals will need to evaluate the cybersecurity measures companies are taking for IoT devices connected to their network to determine risk and evaluate claims.

V. Conclusion

Claims professionals need to be equipped to handle the claims arising out of artificial intelligence and IoT devices that do not function as designed or intended. Claims processing will need to be revamped so that the claims professional has the tools and resources at the beginning of a claim to obtain the necessary information to properly evaluate liability. . Evaluating accidents involving IoT devices and artificial intelligence is unique and requires an understanding of how IoT and AI contributed to accidents. Ongoing education of claims professionals on technology developments and the legal liabilities associated with IoT failures and artificial intelligence unintended consequences will be critical to managing risk.