Framing Public Policy for Adversarial Machine Learning

STS Research Paper
Presented to the Faculty of the
School of Engineering and Applied Science
University of Virginia

By

Jihyeong Lee

April 20th, 2020

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Signed: ____

Approved: _______________________________________ Date ________________________
Rider Foley, Department of Engineering and Society

Introduction

From self-driving cars to fraud detection, algorithms are being widely adopted in many facets of society, requiring stronger guarantees of safety and effectiveness. Machine learning and artificial intelligence (ML & AI) are computer algorithms capable of learning how to complete these tasks automatically, becoming increasingly adopted for a variety of fields. One threat to ML performance and safety is adversarial machine learning, which is the method of causing ML models to behave in an undesirable way through "adversarial attacks". Specifically, an adversarial attack may trick a model into giving an incorrect output, prevent it from learning correctly, or even reveal secure information about itself or the data it was trained on (Liu et al., 2018). For example, a few markings on a stop sign may cause a self-driving car to incorrectly interpret a stop sign as a speed limit sign (Eykholt et al., 2018). As ML and AI become more prevalent, it is important to ensure that models are robust, or in other words, work properly under a wide variety of conditions.

Currently, adversarial attacks pose complex problems from both technical and policy standpoints. From a technical standpoint, building defenses against adversarial attacks has proven difficult for researchers, as many attacks exploit some of the fundamental structures of how ML models are built (Liu et al., 2018). For example, many image tagging algorithms simply look for mathematical similarities between the pixel values of images, which leaves algorithms vulnerable to attacks that exploit the mathematical functions underlying the images. Since these vulnerabilities are built into the fundamental structure of machine learning algorithms, they are difficult to fix (Liu et al., 2018). From a policy standpoint, we can examine the Computer Fraud and Abuse Act (CFAA), the primary legislation that covers cybercrimes in the United States. In response to the lack of specific legislation to address adversarial attacks and general machine learning safety, several researchers have analyzed existing legislation and how it interacts with adversarial attacks. For instance, adversarial attacks could be considered as a case of transmitting unauthorized code to the model, and thus may be covered under the CFAA (Kumar, O'Brien, Albert, & Vilojen, 2018). However, the CFAA is not comprehensive in dealing with the implications of adversarial attacks. For example, liability legislation is still unclear on whether the company or the attacker should be held liable for the damage caused by a model misbehaving due to an attack (Kumar et al., 2018). Thus, machine learning and adversarial attacks will require new legal and regulatory frameworks. With this challenge in mind, this study attempts to answer the question: How will public policy be shaped in order to mitigate the threats posed by adversarial machine learning?

Framing Public Policy Around Adversarial Machine Learning

Currently, there is a lack of regulation for dealing with adversarial machine learning, with the exception being the CFAA, which is one of the primary pieces of legislation that regulates cybercrimes in the United States (Kumar et al. 2018). However, CFAA critics reveal its limitations. In a joint essay by legal scholars and machine learning researchers, scholars argue that the wording of the CFAA is too ambiguous, making it unsuitable for regulating adversarial attacks (Calo, Evtimov, Fernandes, Kohno, & O'Hair, 2018). Through a theoretical analysis of case studies, researchers found several potential attacks that can be interpreted in either direction of being covered or not covered under the CFAA.

Thus, new legal frameworks are required in order to deal with the upcoming issue of adversarial attacks. However, in order to determine the best solution and the set of rules that new regulation or legislation around adversarial attacks should focus on, we must consider how machine learning models are exposed to adversarial threats in the first place and what an appropriate response should be. A study by researchers at Microsoft found that these attacks are becoming easier within our technologically involved society (Marshall, Rojas, Stokes, & Brinkman, 2018). One reason is that ML models are becoming increasingly dependent on public data sources and live interactions, which leaves them vulnerable to malicious attackers messing with training data. Furthermore, the researchers believe that there needs to be a shift in many aspects of deploying models, including security practices and the ability to diagnose these models easier. Therefore, rather than providing strictly technical recommendations for combatting adversarial attacks, the authors believe studying the interaction between the model, its users, its training data, and its developers will lead to better solutions (Marshall et al., 2018).

Two STS frameworks will be effective in exploring this topic: Actor Network Theory (ANT) and Anticipatory Governance (AG). ANT explores the interactions between the technology and society by treating both as equal "actors" in a complex network of interactions (Callon & Blackwell, 2007). Since ANT emphasizes the equal weighting of human and non-human actors in determining the cause and effects of changes, this method is effective in studying how society delegates responsibilities to technology, and in turn, how technology shapes society. As implied by the aforementioned study by the Microsoft researchers, ML systems will depend on interactions with the public, outside the control of its creators, and thus, all interacting parties of the network should be examined carefully (Marshall et al., 2018). As ML becomes more widely adopted across a variety of fields, adversarial attacks to these systems will become more common, requiring us to acknowledge and emphasize the complex network to identify and diagnose threats to the system.

Anticipatory governance studies how to govern technologies before they exist, "anticipating" the problems that might arise from this technology (Guston, 2014). This requires both an analysis of historical impacts of similar technologies, as well as a creative approach to anticipating new problems that come with the new technical aspects. Machine learning and adversarial attacks are still new, meaning we have not encountered many problems in reality. Thus, governments need to take a proactive approach, making anticipatory governance a solid candidate for analyzing the problem and developing new solutions.

Machine Learning and Adversarial Attacks

Machine learning is a technique of building a software program that can infer the relationships between an input and a desired output (Géron, 2019). This is often done by taking a large data set then feeding it into a learning algorithm that will build an approximate mathematical representation of the data. One example of a common machine learning algorithm is the neural network, an algorithm roughly modeled after the human brain - information from an input flowing between multiple layers of artificial neurons. This layered complexity allows neural networks to recognize patterns and learn how to perform tasks similar to humans, such as driving a car or labeling images. Other examples include decision trees, linear and logistic regression, and data clustering, all of which are based on methods from traditional statistics (Géron, 2019).

Due to their approximate nature, machine learning algorithms do not build a perfect representation of the world, which leaves structural flaws in algorithms. Adversarial machine learning entails ways to "trick" a ML algorithm by exploiting these flaws. There are several methods ranging from preventing a model from learning the correct relationships, causing the model to leak data, or forcing the model to behave undesirably (Liu et al., 2018). For example, as shown in Figure 1, Eykholt et al. (2018) present adversarial attacks that could be used to fool image recognition algorithms used to identify road signs, by examining the mathematical structure of a neural network and changing pixel values to force errors. This specific attack is largely theoretical so far, as they require the attacker to perform precise perturbations to objects, making it nearly impossible to conduct in the physical world. Another widely used example is the email spam filter. These systems often deploy pattern recognition machine learning algorithms that can be mathematically exploited to craft new spam emails that can evade the filters, creating a constant battle between email service providers and spammers. Research has even shown that it is theoretically impossible to create a perfect email filter that can catch all spam email, suggesting that adversarial attacks are not issues that can be permanently fixed solely through technical solutions (Tygar, 2011).

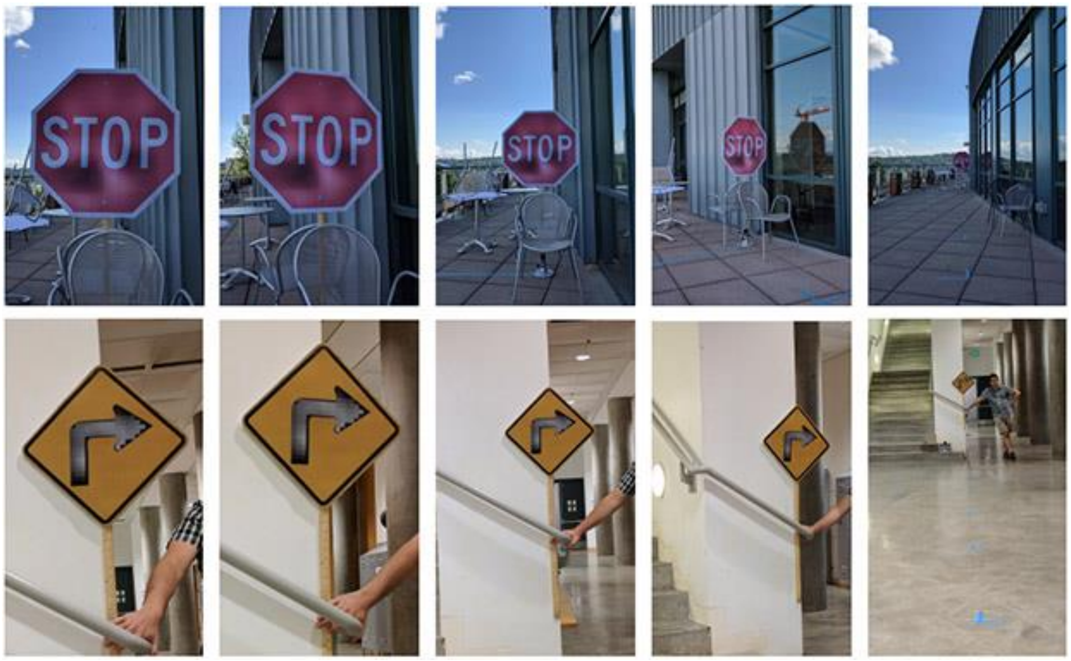

Figure 1. Examples of adversarial attacks on image recognition neural networks. The stop signs are interpreted as 45 MPH speed limit signs and the right turn signs are classified as stop signs (Image source: Eykholt et al, 2018)

Research Question and Methods

The existence of adversarial attacks threatens the future of the trustworthiness, and ultimately, the adoption of machine learning. Protecting against attacks pose many technical challenges, and the lack of public policy so far puts the technology in an even more precarious state. Thus, I explored the following question through my thesis: Is public policy in the United States currently being shaped adequately in order to mitigate the threats posed by adversarial machine learning?

To begin, I expanded upon how the Computer Fraud and Abuse Act (CFAA) can apply to the domain of adversarial machine learning as a method of deterrence. This was primarily done by reviewing the legal essay *Is Tricking a Robot Hacking,* which examines scenarios in which adversarial attacks pose challenges to the application of the CFAA. (Calo, Evtimov, Fernandes, Kohno, & O'Hair, 2018). Next, I examined the policies two governments placed machine learning on their public policy agendas. The first is the United States, which has begun its "America AI Initiative" under Executive Order No. 13859 (E.O. 13859). The executive order sets AI R&D at the forefront of national priorities, and places explicit focus on "making AI trustworthy". The executive order has spun off multiple policy documents on monitoring the development of AI/ML in the United States, such as a plan to guide the research and development (R&D) progress into AI (Artificial Intelligence R&D Interagency Working Group [AI R&D IWG], 2019). Additionally, the National Institute of Standards and Technology (NIST) plans to develop standards for AI/ML, considering several aspects including security and safety (National Institute of Standards and Technology [NIST], 2019). Another important document is a set of guidelines on regulating AI, which gives direct insight into how the federal agencies plan on addressing (or not addressing) adversarial attacks and AI security (Office of Management and Budge [OMB], 2019).

The second is the European Union (EU) and its Guidelines for Ethical AI (High-Level Expert Group on AI [AI HLEG], 2019a). This is a framework released by the EU on how engineers should build safe and ethical ML systems. The guidelines dedicate an entire section on providing a checklist for building a robust system against adversarial attacks (AI HLEG, 2019a). This set of guidelines are currently undergoing live testing with private corporations, but results have not been released as of writing this thesis. Furthermore, the EU has released policy and investment recommendations based on its ethical framework, showing a preview of the policy changes it plans to make in the near future (AI HLEG, 2019b).

These governments provide two different social contexts for regulating AI/ML: the EU has already shown willingness to heavily regulate technologies through the GDPR, while the US has been more conservative in enacting new laws and regulations. This may be reflected in the specificity and restrictiveness of government standards and regulations of AI/ML safety, with the EU potentially placing more restrictions on companies using AI/ML than the US. From the AG framework, we can see that both governments are willing to tackle the issue proactively, despite adversarial attacks being largely hypothetical threats to AI. From the ANT framework, these policy guidelines seem to consider the full network of interactions between an AI system, its users, creators, and attackers, with those considerations being directly translated into the policy itself.

Results

The United States is still in its early stages of addressing the upcoming threats, but still placing emphasis on the importance of this issue. The newly proposed set of frameworks under the American AI Initiative have been focused on the need for practicing proper risk management by emphasizing resilience and robustness in system development, keeping the threats in mind during the design phase. On the other hand, the CFAA has been a controversial piece of legislation due to its vagueness, and thus may be too vague to deal with adversarial attacks, as discussed below.

Examination of the Computer Fraud and Abuse Act of 1984 (CFAA)

The effectiveness of CFAA has faced much criticism for how expansive and vague it has been in its wording. The CFAA's vagueness has forced the burden of identifying what specific types of hacking constitutes as cybercrime under the law on the courts, rather than lawmakers (Kerr, 2009). Simmons (2016) similarly argues that the CFAA is too vague as well, since the legislation was originally developed when computer systems were simpler. The original goal of the CFAA was to cover computer-related crimes in cases where existing law failed to do so. For example, trespassing is a concept that's easy to define in the physical world. However, for a computer user, the definition is fuzzier: Is an employee accessing a dataset for some purpose other than work a case of trespassing? What about a Terms of Service violation? Should that be a criminal or civil charge? These are the types questions that the CFAA was supposed to answer. However, as computer technology evolved, the definitions of the CFAA were often too vague or required constant amendment. Simmons discusses *United States v. Lowson,* a case were the defendants wrote a script to quickly purchase concert tickets by bypassing a website's antiscripting measures, which led to criminal charges. Simmons (2016) argues that the same act would not have been a crime if the defendants had done the same action by hand, creating an ambiguous boundary between physical and cyber-crimes.

This fact could apply in both directions when it comes to adversarial attacks: the courts could either expand its definitions of what constitutes as a cybercrime, or determine that adversarial attacks are too different in nature. For example, consider the Department of Justice's Office of Legal Education (2010) prosecution guidelines for the CFAA, which states that "knowingly" transmitting information without access to a protected computer to cause damage constitutes as a misdemeanor. Thus, on the surface level, it seems that current law provides coverage to deal with adversarial attacks.

On a deeper level, the situation becomes much more ambiguous. Through a legal essay, Calo et al. (2018) extensively examines some of the issues of the CFAA's incompatibility with regulating adversarial attacks. One hypothetical scenario the essay provides is the act of purchasing a TV ad to embed an adversarial attack in the audio to manipulate a voice-activated home assistant system. The adversarial attack in this case was conducted against an individual, which means that the device is not a "protected computer", i.e. a computer involved in interstate commerce, the financial sector, or the United States government. In a different example, users partaking in surveys to collect data for a credit rating system cause the system to learn a spurious relationship between skateboarding and good credit. While this may seem like it caused damage to a financial system, a previous ruling has shown that this access must have bypassed a security protocol that forbade such behavior, so this case is unclear as well (*United States v. Kane*, 2013).

Finally, security researchers have criticized the CFAA for not excluding "white hat" hacking, i.e. hacking to find vulnerabilities so that companies can fix them (Etovich and Van der

Merwe). In a constantly evolving field such as cybersecurity, it is important to incentivize “white hat” hacking, as it is impossible to create a perfectly secure software system. The same applies to adversarial attacks, where users could report vulnerabilities of systems to certain attacks.

Proposed Frameworks for Regulating Artificial Intelligence

As opposed to the CFAA, which dealt with cybercrimes through deterrence of crime, the US and EU have begun planning frameworks for ensuring that AI is used for the benefit of the country. Both the United States and the European Union have focused on the goal of establishing “trustworthiness” with AI systems, and have mostly adopted similar approaches. In the United States, these values have been communicated to the public through a variety of White House memorandums under the “American AI Initiative”. In the European Union, the High-Level Ethics Group on AI (AI HLEG) put-forth ethical guidelines on how AI should be designed and used in the EU, as well as accompanying policy recommendations for the guidelines. Relevant information regarding adversarial attacks from both plans are summarized in Table 1.

The main takeaway from these documents is that both governments have explicitly considered the dangers of adversarial attacks on future applications of AI, emphasizing the importance of technical robustness, security, and explainability (OMB 2019, AI HLEG 2019a). One difference is that although both governments recognize that heavy regulation could stifle innovation, the American AI Initiative places non-regulation as an explicit strategy in its policy, whereas the EU does not. This is reflected by the US government’s heavier emphasis on building defenses through improvement of the technology, rather than providing more detailed regulatory guidelines like the EU (OMB, 2019; AI HLEG, 2019b). Furthermore, the US leaves the task of regulation to executive agencies and departments. Rather than focusing on generalized guidelines, the American AI Initiative recognizes that the impact of attacks will depend on the application and the users of the AI system itself, thus leaving the task of regulation to the specific application field. Lastly, both governments emphasize the importance of public participation, which is a large component of effective anticipatory governance (OMB, 2019; AI HLEG, 2019b). For instance, the American AI Initiatives invites the public to “participate in all stages of the rulemaking process”, utilizing the power of the public to come up with wider possibilities for scenario planning that will hopefully help inform policymakers (OMB, 2019).

	Both	United States	European Union
Core Values and Execution Plan	• Proposed guidelines and recommendations over actual policy so far	• Executive order 13859: American AI Initiative with 5 pillars: (1) invest in AI research and development (2) unleash AI resources (3) remove barriers to AI innovation (4) train an AI-ready workforce (5) promote an international environment that is supportive of American AI innovation and its responsible use (E.O 13859, 2019) • Calls upon federal agencies and departments to execute the necessary steps • Proposed guidelines for AI regulatory principles (OMB 2019)	• Creating "Trustworthy AI" through 3 main components of how AI should be created and used: (1) Lawful AI (2) Ethical AI (3) Robust AI (AI HLEG, 2019a) • Ethics Guidelines document focusing on 7 Key Requirements, including robustness and safety (AI HLEG, 2019a) • Policy Recommendation document, which complements the Ethics Guidelines by stating policy and investment goals for hitting the 7 Key Requirements (AI HLEG, 2019b)
Focus on Adversarial Attacks and General Safety	• Emphasis on technical robustness, security, and explainability (OMB 2019, AI HLEG, 2019a)	• Heavy emphasis on building defenses by R&D investments into countering attacks (AI R&D IWG, 2019)	• Provides a developer's checklist for developing a secure AI system (AI HLEG, 2019a)
Policy Recommendations	• Emphasizes risk management on the organizational level, depending on the criticality of the AI system (OMB 2019, AI HLEG, 2019a) • Proposes the creation of technical standards for AI security (NIST, 2019; AI HLEG, 2019b) • Avoid heavy regulation that could stifle innovation (OMB, 2019; AI HLEG, 2019b) • Public participation in creating regulation and considering the	• Direct government investments into R&D to develop technical safety, such as research grants (AI R&D IWG, 2019) • More emphasis on avoiding unnecessary regulation through longer explanation of why overregulation could be detrimental to the adoption of AI (OMB, 2019)	• Recommends review of existing legislation and regulations for alignment/misalignment with the ethics guidelines and dealing with the novelty of AI (AI HLEG, 2019b) • Development of "AI-Specific Cybersecurity Infrastructures" (AI HLEG, 2019b) • Mandatory requirements for critical AI systems to "conduct a trustworthy AI assessment" and enforce auditability in case of failure (AI HLEG, 2019b)

Table 1. Summary of documents released under the American AI Initiative and the European Union High-Level Ethics Group on AI reports (Lee, 2020).

Discussion

The American AI Initiative provides a more generalized approach than the CFAA by establishing a set of guidelines that governmental agencies, AI developers, and various stakeholders should take. Furthermore, similar to the NIST Cybersecurity Standards, the Initiative aims to provide a set of technical standards that can be used to not only bolster AI safety across all applications, but also provide a set of legal standards that can be used in court (NIST 2018).

Furthermore, through the Initiative, the US government respects the relationship between the technology, regulation, and human interactions. In the lens of the ANT framework, the American AI Initiative views the interaction between malicious users and the technology as part of a larger network of researchers, human stakeholders such as investors and users, and regulatory environments that evolves together. This contrasts targeted legislation such as the CFAA, which simply focuses the relationship between hackers and the owners of a computer system. We can see this consideration from the emphasis on conducting cost-benefit analysis for any regulation in order to derive benefit from the technology without exposing users to danger or stifling researchers from pursuing advancements in AI. This follows the practices of effective anticipatory governance, which asserts that the potential impact of new technologies is never predictable, and the best approach is through a balance of generalized risk management and scenario planning. In this case, we can see that the US respects the threat of adversarial attacks on AI, but takes a cautious approach against placing overly restrictive regulations in consideration of the research community and the adopters of AI.

Unfortunately, there has been no mention of reforming the CFAA or creating new punitive legislation against adversarial attacks, which is worrying considering the shortcomings of the CFAA. However, it is also possible that the Initiative aims to replace dated legislation with standards similar to the aforementioned NIST Cybersecurity Standards, although there is no direct evidence of this in the documents. This may be due to the fact that adversarial attacks and AI itself are still emergent technologies, and thus do not warrant complete legislation yet, but rather, flexible policy guidelines that can instead lay out core principles for faster moving regulation. In the lens of the Anticipatory Governance framework, both governments are still in the planning phase, and looking to develop insights into potential scenarios of how threats to AI might appear. This includes the emphasis on public participation. The American AI Initiatives even invites the public to "participate in all stages of the rulemaking process", utilizing the power of the public to come up with wider scenarios for scenario planning. This fits in with the AG framework, which places importance on utilizing the creativity of the public in identifying threats that a smaller group of experts could miss.

Limitations

This study was limited in several ways. First, most of the analysis on the CFAA and its extension to adversarial attacks has been based on speculation of legal experts, as opposed to actual legal precedence that forms the basis of common law, hence why the Actor Network Theory framework needed to be complemented by an analytical framework that is more future looking, such as the Anticipatory Governance framework. The information is based off of a snapshot of new efforts to regulate a completely novel technology, and thus, is an incomplete view of the government's full plans, and is also susceptible to change. Furthermore, it lacks the perspectives of how non-governmental stakeholders would view these plans, which could have been collected through interviews with experts. Additionally, this study only examined the proposed policies of the US and the EU, as not many other countries have begun creating regulation in controlling AI. However, as the technology develops, more and more governments will need to eventually address this issue, which can be used for further analysis in the future.

Conclusion

Through the analysis of the CFAA and the proposed policy documents, it is evident that the US is taking a more anticipatory approach to addressing this new type of cyber-threat, and also considering the entire network of relationships that are formed around complex technologies. While this is something to be commended, one concern is the lack of attention into the CFAA, which still remains the only punitive legislation for addressing cybercrime in the US. This study shows not only the benefits of taking an anticipatory approach to governance, but also some of its shortcoming, as we tend to ignore the present when we become too focused into managing the future. My hope is that through this study, engineers become aware of these technological issues, and take advantage of the government's plans for public participation and keep the government both forward looking and reflective of the present. Future studies could potentially examine the opinions of technical experts in the field, and whether focusing on building technical defenses is a feasible task. Furthermore, as more countries begin to adopt AI, it is important to examine how different cultures and political environments lead to different policies, and derive policy ideas that we could potentially adopt ourselves. Finally, future studies could also examine the broader topic of risk management for the adoption of AI itself, and whether the benefits are worth the all of the potential threats, including adversarial attacks.

References

Artificial Intelligence R&D Interagency Working Group. (2019). *2016–2019 Progress Report: Advancing Artificial Intelligence R&D*. Retrieved from https://www.whitehouse.gov/wp-content/uploads/2019/11/AI-Research-and-Development-Progress-Report-2016-2019.pdf

Callon, M., & Blackwell, O. (2007). Actor-network theory. *The Politics of Interventions, Oslo Academic Press, Unipub, Oslo*, 273-286.

Calo, R., Evtimov, I., Fernandes, E., Kohno, T., & O'Hair, D. (2018). Is Tricking a Robot Hacking?. *University of Washington School of Law Research Paper*, (2018-05).

Etcovich, D., & Van der Merwe, T. (2018). *Coming in from the Cold: A Safe Harbor from the CFAA and the DMCA §1201 for Security Researchers* (Research Publication No. No. 2018-4). Berkman Klein Center for Internet & Society at Harvard University. Retrieved from https://cyber.harvard.edu/publication/2018/coming-cold-safe-harbor-cfaa-and-dmca-ss1201

Exec. Order No. 13,859, 3 C.F.R. 84 Fed. Reg. 3967 (February 11, 2019).

Eykholt, K., Evtimov, I., Fernandes, E., Li, B., Rahmati, A., Xiao, C., ..., Song, D. (2018). Robust Physical-World Attacks on Deep Learning Models. *Conference on Computer Vision and Pattern Recognition.*

Géron Aurélien. (2019). *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: concepts, tools, and techniques to build intelligent systems*. Sebastopol, CA: OReilly Media, Inc.

Guston, D. H. (2014). Understanding ‘Anticipatory Governance’. *Social Studies of Science, 44*(2), 218-242.

High-Level Expert Group on AI (2019a). Ethics Guidelines for Trustworthy AI. Retrieved from https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

High-Level Expert Group on AI (2019b). Policy and investment recommendations for trustworthy Artificial Intelligence. Retrieved from https://ec.europa.eu/digital-single-market/en/news/policy-and-investment-recommendations-trustworthy-artificial-intelligence

Kerr, O. (2009). Vagueness Challenges to the Computer Fraud and Abuse Act. *Minnesota Law Review*, *94*, 1561–1587. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1527187.

Kumar, R. S. S., O'Brien, D. R., Albert, K., & Vilojen, S. (2018). *Law and Adversarial Machine Learning*. *arXiv Preprint arXiv:1810.10731*.

Lin, Y. C., Hong, Z. W., Liao, Y. H., Shih, M. L., Liu, M. Y., & Sun, M. 2017). Tactics of Adversarial Attack on Deep Reinforcement Learning Agents. *International Joint Conference on Artificial Intelligence*. Melbourne, Australia.

Liu, Q., Li, P., Zhao, W., Cai, W., Yu, S., & Leung, V. C. (2018). A Survey on Security Threats and Defensive Techniques of Machine Learning: A data driven view. *IEEE Access*, 6, 12103-12117.

Marshall, A., Rojas, R., Stokes, J., & Brinkman, D. (2018, December 2). Securing the Future of AI and ML at Microsoft - Security Documentation. Retrieved October 17, 2019, from https://docs.microsoft.com/en-us/security/securing-artificial-intelligence-machine-learning.

National Institute for Standards and Technology. (2018). Uses and Benefits of the Framework. Retrieved from https://www.nist.gov/cyberframework/online-learning/uses-and-benefits-framework

National Institute for Standards and Technology. (2019). *U.S. Leadership in AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools*. U.S. Department of Commerce. Retrieved from https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

Office of Legal Education. (2010). *Prosecuting Computer Crimes* (Manual). The United States Department of Justice. Retrieved from https://www.justice.gov/sites/default/files/criminal-ccips/legacy/2015/01/14/ccmanual.pdf

Office of Management and Budget. (2019). *Guidance for Regulation of Artificial Intelligence Applications* (Draft Memorandum No. 85 FR 1825). Executive Office of the President. Retrieved from https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-of-AI-1-7-19.pdf

Shen, L. (2014). The NIST Cybersecurity Framework: Overview and Potential Impacts. *Scitech Lawyer*, 10(4), 16-19.

Simmons, R. (2016). The Failure of the Computer Fraud and Abuse Act: Time to Take an Administrative Approach to Regulating Computer Crime. *George Washington Law Review*, 84, 1703-1724.

Tygar, J. D. “Adversarial Machine Learning.” *IEEE Internet Computing*, vol. 15, no. 5, Sept. 2011, pp. 4–6. *DOI.org (Crossref)*, doi:10.1109/MIC.2011.112.

United States v. Kane, No. 11-mj-0000 (United States District Court for the District of Nevada 2013).