

Applied Artificial Intelligence in Modern Warfare & National Security Policy

Brian Seamus Haney¹

Abstract

Artificial Intelligence (AI) applications in modern warfare have revolutionized national security power dynamics between the United States, China, Russia, and private industry. The United States has fallen behind in military technologies and is now at the mercy of big technology companies to maintain peace. After committing \$150 billion toward the goal of becoming the AI technology world leader, China claimed success in 2018. In 2019, Chinese researchers published open-source code for AI missile systems controlled by deep reinforcement learning algorithms. Further, Russia's continued interference in U.S. elections has largely been driven by AI applications in cybersecurity. Yet, despite outspending Russia and China combined on defense, the United States is failing to keep pace with foreign adversaries in the AI arms race.

Previous legal scholarship dismisses AI militarization as futuristic science-fiction, accepting without support the United States' prominence as the world leader in military technology. This inter-disciplinary Article provides three main contributions to legal scholarship. First, this is the first piece in legal scholarship to take an informatics-based approach toward analyzing the range of AI applications in modern warfare. Second, this is the first piece in legal scholarship to take an informatics-based approach in analyzing national security policy. Third, this is the first piece to explore the complex power and security dynamics between the United States, China, Russia, and private corporations in the AI arms race. Ultimately, a new era of advanced weaponry is developing, and the United States Government is sitting on the sidelines.

¹ J.D. Notre Dame Law School 2018, B.A. Washington & Jefferson College 2015. Special thanks to Richard Susskind, Margaret Cuonzo, Max Tegmark, Ethem Alpaydin, Sam Altman, Josh Achiam, Volodymyr Mnih, and Angela Elias.

Table of Contents

<i>INTRODUCTION</i>	3
<i>I. ARTIFICIAL INTELLIGENCE</i>	5
A. DEEP LEARNING	6
B. REINFORCEMENT LEARNING	10
C. DEEP REINFORCEMENT LEARNING	12
<i>II. SECURITY THREATS</i>	19
A. MISSILES.....	20
B. CYBERSECURITY	24
C. ARTIFICIAL GENERAL INTELLIGENCE.....	30
<i>III. POLICY</i>	32
A. CHINA	34
B. RUSSIA	36
C. UNITED STATES	37
<i>CONCLUSION</i>	43
<i>APPENDIX A. SUMMARY OF NOTATION</i>	46
<i>APPENDIX B. SIGNIFICANT CYBER INCIDENTS DATA</i>	47

INTRODUCTION

Cyberwarfare continues on a daily basis, with the United States under constant attack.² Threats of nuclear missile strikes from adversaries appear in daily headlines.³ Today, Artificial Intelligence (AI) is the United States' most powerful weapon for defense.⁴ Yet, in AI development, the United States is falling behind adversaries like China and Russia.⁵ In 2017, China committed \$150 billion toward becoming the world leader in AI, claiming success the next year.⁶ Interference in U.S. elections is largely being driven by substantial Russian investments in AI cybersecurity applications.⁷ All the while, the United States Government and Department of Defense remain at the mercy of big technology companies like Google and Microsoft to ensure advancements in AI research and development.⁸

The Law of Accelerating Returns (“LOAR”) states that fundamental measures of information technology follow predictable and exponential trajectories.⁹ Indeed, information

² James Kadtko, John Wharton, *Technology and National Security: The United States at a Critical Crossroads*, Defense Horizons, National Defense University at 1 (March 2018).

³ Hyung-Jin Kim, *North Korea Confirms Second Test of a Multiple Rocket Launcher*, Time (September 11, 2019), <https://time.com/5673813/north-korea-confirms-second-rocket-launcher-test/>. *See also* Nasser Karimi, John Gambrell, *Iran uses Advanced Centrifuges, Threatens Higher Enrichment*, Associated Press (September 7, 2019), <https://www.apnews.com/7e896f8a1b0c40769b54ed4f98a0f5e6>.

⁴ Karl Manheim, Lyric Kaplan, *Artificial Intelligence: Risks to Privacy and Democracy*, 21 YALE J. L. & TECH. 106, 108 (2019). *See also* *Elon Musk at the National Governors Association 2017 Summer Meeting*, C-SPAN, 50:00 (July 15, 2017) <https://www.c-span.org/video/?c4676772/elon-musk-national-governors-association-2017-summer-meeting>.

⁵ Kadtko, Wharton, *supra* note 2, at 1.

⁶ Gregory C. Allen, *Understanding China's AI Strategy: Clues to Chinese Strategic Thinking on Artificial Intelligence and National Security*, Center for a New American Security 9 (February 2019).

⁷ Congressional Research Service, *Artificial Intelligence and National Security* 24 (January 30, 2019), <https://fas.org/sgp/crs/natsec/R45178.pdf>.

⁸ Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 354 (2016).

⁹ RAY KURZWEIL, *HOW TO CREATE A MIND* 250 (2012).

technologies build on themselves in an exponential manner.¹⁰ Applied to AI, the LOAR provides strong support for AI's increasing role in protecting the national defense.¹¹ Indeed, similar to the way in which aviation and nuclear weapons transformed the military landscape in the twentieth century, AI is reconstructing the fundamental nature of military technologies today.¹²

Yet legal scholars continue to deny and ignore AI's applications as a weapon of mass destruction. For example, in a recent MIT Starr Forum Report, the Honorable James E Baker, former Chief Judge of the United States Court of Appeals for the Armed Forces, argues "we really won't need to worry about the long-term existential risks."¹³ And, University of Washington Law Professor, Ryan Calo argues, regulators should not be distracted by claims of an "AI Apocalypse" and to focus their efforts on "more immediate harms."¹⁴ All the while, private corporations are pouring billions into AI research, development, and deployment.¹⁵ In a 2019 interview, Paul M. Nakasone, The Director of the National Security Agency (NSA) stated, "I suspect that AI will play a future role in helping us discern vulnerabilities quicker and allow us to focus on options that will have a higher likelihood of success."¹⁶ Yet, Elon Musk argues today, "The biggest risk that we face as a civilization is artificial intelligence."¹⁷ The variance in the position of industry leaders relating to AI and defense demonstrates a glaring disconnect and information gap between legal scholars, Government leaders, and private industry.

¹⁰ *Id.* at 256.

¹¹ NICK BOSTROM, *SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES* 94 (2017).

¹² Honorable James E Baker, *Artificial Intelligence and National Security Law: A Dangerous Nonchalance*, Starr Forum Report 1 (2018).

¹³ *Id.* at 5.

¹⁴ Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399, 431 (2017).

¹⁵ Andrew Thompson, *The Committee on Foreign Investment in The United States: An Analysis of the Foreign Investment Risk Review Modernization Act of 2018*, 19 J. HIGH TECH. L. 361, 363 (2019).

¹⁶ An Interview with Paul M. Nakasone, National Defense University Press, JFQ 92, 1st Quarter, at 9 (2019).

¹⁷ *Elon Musk at the National Governors Association 2017 Summer Meeting*, C-SPAN, 50:00 (July 15, 2017) <https://www.c-span.org/video/?c4676772/elon-musk-national-governors-association-2017-summer-meeting>.

The purpose of this Article is to aid in closing the information gap, explaining the applications of AI in modern warfare. Further, this Article contributes the first informatics-based analysis of the national security policy landscape. This Article proceeds in three parts. Part I explains the state-of-the-art in AI technology. Part II explores three national security threats resulting from AI applications in modern warfare. Part III discusses national security policy relating to AI from international and domestic perspectives.

I. ARTIFICIAL INTELLIGENCE

Contemporary scholars have presented several different definitions of AI. For example, MIT Professor Max Tegmark concisely defines intelligence as the ability to achieve goals¹⁸ and AI as “non-biological intelligence.”¹⁹ Additionally, according to Stanford Professor Nils Nilsson AI is “concerned with intelligent behavior in artifacts.”²⁰ A recent One Hundred Years Study defines AI as, “a science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their nervous systems and bodies to sense, learn, reason, and take action.”²¹ For the purposes of this paper AI is any system replicating the thoughtful processes associated with human thought.²² Advancements in AI technologies continue at alarming rates.²³ This Part proceeds by discussing three types of AI systems commonly used in the context of national security: deep learning, reinforcement learning, and deep reinforcement learning.

¹⁸ MAX TEGMARK, LIFE 3.0 BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE 50 (2017).

¹⁹ *Id.* at 39.

²⁰ NILS J. NILSSON, ARTIFICIAL INTELLIGENCE: A NEW SYNTHESIS 1 (1998).

²¹ Artificial Intelligence and Life in 2030, One Hundred Year Study on Artificial Intelligence, Stanford University 1 (2016).

²² Brian S. Haney, *The Perils and Promises of Artificial General Intelligence*, 45 J. Legis. 151, 152 (2018).

²³ PAUL E. CERUZZI, COMPUTING A CONCISE HISTORY 114 (2012).

A. DEEP LEARNING

Deep learning is a process by which neural networks learn from large amounts of data.²⁴ Defined, data is any recorded information about the world.²⁵ In deep learning, the idea is to learn feature levels of increasing abstraction with minimum human contribution.²⁶ The models inspiring current deep learning architectures have been around since the 1950s.²⁷ Indeed, the Perceptron, which serves as the basic tool of neural networks was proposed by Frank Rosenblatt in 1957.²⁸ However, artificial intelligence research remained relatively unprosperous until the dawn of the internet.²⁹ Generally, deep learning systems are developed in four parts: data pre-processing, model design, training, and testing.

Deep learning is all about the data.³⁰ Every two days humans create more data than the total amount of data created from the dawn of humanity until 2003.³¹ Indeed, the internet is the driving force behind modern deep learning strategies because the internet has enabled humanity to organize and aggregate massive amounts of data.³² According to machine learning scholar, Ethem Alpaydin, it's the data that drives the operation, not human programmers.³³ The majority of the time spent with deep learning system development is during the pre-processing stage.³⁴ During this initial phase, machine learning researchers gather, organize, and aggregate data to be analyzed by neural networks.³⁵

²⁴ Haney, *supra* note 22 at 157.

²⁵ ETHEM ALPAYDIN, *MACHINE LEARNING* 3 (2016). *See also* MICHAEL BUCKLAND, *INFORMATION AND SOCIETY* 21-22 (2017). (Discussing definitions of information.)

²⁶ JOHN D. KELLEHER, BRENDEN TIERNEY, *DATA SCIENCE* 134 (2018).

²⁷ SEBASTIAN RASCHKA, VAHID MIRJALILI, *PYTHON MACHINE LEARNING* 18 (2017).

²⁸ *Id.*

²⁹ PETER J. DENNING, MATTI TEDRE, *COMPUTATIONAL THINKING* 93 (2019).

³⁰ David Lehr, Paul Ohm, *Playing with The Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 668 (2017).

³¹ RICHARD SUSSKIND, *TOMORROW'S LAWYERS* 11 (2017).

³² ALPAYDIN, *supra* note 25, at 10-11.

³³ *Id.* at 12.

³⁴ KELLEHER, TIERNEY, *supra* note 26, at 97.

³⁵ *Id.*

The types of data neural networks process vary.³⁶ In the context of autonomous warfare systems, one example may be images stored as pixel values to be associated with object classification for targeting.³⁷ Another example is gaining political insight with a dataset of publicly available personal data on foreign officials. How the data is organized is in large part dependent on the goal of the deep learning system.³⁸ If a system is being developed for predictive purposes the data may be labeled with positive and negative instances of an occurrence.³⁹ Or, if the system is being learned to gain insight, the data may remain unstructured, allowing the model to complete the organization task.⁴⁰

A deep learning system's model is the part of the system which analyzes the information.⁴¹ Most commonly the model is a neural network.⁴² Neural networks serve the function of associating information to derive knowledge.⁴³ Neural networks are modeled based upon the biological neo-cortex.⁴⁴ Indeed, the human brain is composed of processing units called neurons.⁴⁵ Each neuron in the brain is connected to other neurons through structures called synapses.⁴⁶ A biological neuron consists of dendrites—receivers of various electrical impulses from other neurons—that are gathered in the cell body of a neuron.⁴⁷ Once the neuron's cell body has collected enough electrical energy to exceed a threshold amount, the neuron transmits an electrical charge to other neurons in

³⁶ *Id.* at 101.

³⁷ Richard Wu, et al., A Framework Using Machine Vision and Deep Reinforcement Learning for Self-Learning Moving Objects in a Virtual Environment, AAI 2017 Fall Symposium Series (2017), <https://aaai.org/ocs/index.php/FSS/FSS17/paper/view/16003/15319>.

³⁸ Michael Simon, et. al., *Lola v. Skadden and the Automation of the Legal Profession*, 20 Yale J.L. & Tech 234, 253 (2018).

³⁹ TARIQ RASHID, MAKE YOUR OWN NEURAL NETWORK 13 (2018).

⁴⁰ ALPAYDIN, *supra* note 25, at 111.

⁴¹ KELLEHER, TIERNEY, *supra* note 26, at 121.

⁴² TEGMARK, *supra* note 18, at 76.

⁴³ ALPAYDIN, *supra* note 25, at 106-107.

⁴⁴ Michael Simon, et. al., *supra* note 38, 254.

⁴⁵ MOHEB COSTANDI, NEUROPLASTICITY 6 (2016).

⁴⁶ *Id.* at 9.

⁴⁷ *Id.* at 7.

the brain through synapses.⁴⁸ This transfer of information in the biological brain provides the foundation for the way in which modern neural networks operate.⁴⁹

Every neural network has an input layer and an output layer.⁵⁰ However, in between the input and output layer, neural networks contain multiple hidden layers of connected neurons.⁵¹ In a neural network, the neurons are connected by weight coefficients modeling the strength of synapses in the biological brain.⁵² The depth of the network is in large part a description of the number of hidden layers.⁵³ Deep Neural Networks start from raw input and then each hidden layer combines the values in its preceding layer and learns more complicated functions of the input.⁵⁴ The mathematics of the network transferring information from input to output varies, but is generally matrix mathematics and vector calculus.⁵⁵ During training, the model processes data from input to output, often described as the feedforward portion.⁵⁶ The output of the model is typically a prediction.⁵⁷ For example, whether an object is the correct target, or the wrong target would be calculated with a convolutional neural network (CNN).⁵⁸ The function of the CNN is in essence a classification task, where the CNN classifies objects or areas based upon their similarity.⁵⁹ CNNs are the main model used for deep learning in computer vision tasks.⁶⁰

⁴⁸ RASCHKA, MIRJALILI, *supra* note 27, at 18.

⁴⁹ Haney, *supra* note 22 at 158.

⁵⁰ KURZWEIL, *supra* note 9, at 132.

⁵¹ ALPAYDIN, *supra* note 25, at 100.

⁵² *Id.* at 88.

⁵³ TEGMARK, *supra* note 18, at 76.

⁵⁴ ALPAYDIN, *supra* note 25, at 104.

⁵⁵ Manon Legrand, Deep Reinforcement Learning for Autonomous Vehicle Control Among Human Drivers, at 23 (academic year 2016–17) (unpublished C.S. thesis, Université Libre de Bruxelles), https://ai.vub.ac.be/sites/default/files/thesis_legrand.pdf.

⁵⁶ EUGENE CHARNIAK, INTRODUCTION TO DEEP LEARNING 10 (2018).

⁵⁷ Harry Surden, *Machine Learning and Law*, 89 Wash. L. Rev. 87, 90 (2014).

⁵⁸ Daniel Maturana, Sebastian Scherer, 3D Convolutional Neural Networks for Landing Zone Detection from LiDar, 2 (2015), <https://ieeexplore.ieee.org/document/7139679>.

⁵⁹ RASHID, *supra* note 39, at 159.

⁶⁰ Legrand, *supra* note 55, at 23.

However, the learning occurs during the backpropagation process.⁶¹ Backpropagation describes the way which neural networks are trained to derive meaning from data.⁶² Generally, the mathematics of the backpropagation algorithm includes partial derivative calculations and a loss function to be minimized.⁶³ The algorithm's essential function adjusts the weights of a neural network to reduce error.⁶⁴ The algorithm's ultimate goal is the convergence of an optimal network, but probabilistic maximization also provides state-of-the-art performance in real world domains.⁶⁵ Dynamic feedback allows derivative calculations supporting error minimization.⁶⁶ One popular algorithm for backpropagation is stochastic gradient descent (SGD), iteratively updates the weights of the network according to a loss function.⁶⁷

After the training process the model is then tested on new data, and if successful, deployed for the purpose deriving knowledge from information.⁶⁸ The process of deriving knowledge from information is commonly accomplished with feature extraction.⁶⁹ Feature extraction is a method of dimensionality reduction allowing raw inputs to convert to an output revealing abstract relationships among data.⁷⁰ Neural networks extract these abstract relationships by combining previous input information in higher dimensional space as the network iterates.⁷¹ In other words, deep neural networks learn more complicated functions of their initial input when each hidden layer combines the values of the preceding layer.⁷² In addition to deep learning, reinforcement

⁶¹ KELLEHER, TIERNEY, *supra* note 26, at 130.

⁶² ALPAYDIN, *supra* note 25, at 100.

⁶³ PAUL JOHN WERBOS, THE ROOTS OF BACKPROPAGATION FROM ORDERED DERIVATIVES TO NEURAL NETWORKS AND POLITICAL FORECASTING 269 (1994).

⁶⁴ ALPAYDIN, *supra* note 25, at 89.

⁶⁵ KELLEHER, TIERNEY, *supra* note 26, at 131.

⁶⁶ WERBOS, *supra* note 63, at 72.

⁶⁷ Steven M Bellovin, et al., *Privacy and Synthetic Datasets*, 22 STAN. TECH. L. REV. 1, 29 (2019).

⁶⁸ ALPAYDIN, *supra* note 25, at 106-107.

⁶⁹ *Id.* at 89.

⁷⁰ *Id.* at 102.

⁷¹ KELLEHER, TIERNEY, *supra* note 26, at 135.

⁷² ALPAYDIN, *supra* note 25, at 104.

learning is also a major cause of concern for purposes of national security policy.


B. REINFORCEMENT LEARNING

At its core, reinforcement learning is an optimization algorithm.⁷³ In short, reinforcement learning is a type of machine learning concerned with learning how an agent should behave in an environment to maximize a reward.⁷⁴ Agents are the software programs making intelligent decisions.⁷⁵ Generally, reinforcement learning algorithms contain three elements:

1. Model: the description of the agent-environment relationship;
2. Policy: the way in which the agent makes decisions; and
3. Reward: the agent's goal.⁷⁶

The fundamental reinforcement learning model is the Markov Decision Process (MDP).⁷⁷ The MDP model was developed by the Russian Mathematician Andrey Markov in 1913.⁷⁸

Interestingly, Markov's work over a century ago remains the state-of-the-art in AI today.⁷⁹ The model below describes the agent-environment interaction in an MDP:⁸⁰



⁷³ Volodymyr Mnih et al., Human-Level Control Through Deep Reinforcement Learning, 518 NATURE INT'L J. SCI. 529, 529 (2015).

⁷⁴ ALPAYDIN, *supra* note 25, at 127.

⁷⁵ RICHARD S. SUTTON, ANDREW G. BARTO, REINFORCEMENT LEARNING: AN INTRODUCTION 3 (2017).

⁷⁶ Katerina Fragkiadaki, Deep Q Learning, Carnegie Mellon Computer Science, CMU 10703 (Fall 2018), https://www.cs.cmu.edu/~katief/DeepRLFall2018/lecture_DQL_katief2018.pdf.

⁷⁷ Haney, *supra* note 22 at 161.

⁷⁸ Gely P. Basharin, et. al, The Life and Work of A.A. Markov, Linear Algebra and its Applications, Vol. 386, at 15 (2004).

⁷⁹ GEORGE GILDER, LIFE AFTER GOOGLE 75 (2018).

⁸⁰ SUTTON, BARTO, *supra* note 75, at 38. (Model created by author based on illustration at the preceding citation).

The environment is made up of states for each point in time in which the environment exists.⁸¹ The learning begins when the agent takes an initial action selected from the first state in the environment.⁸² Once the agent selects an action, the environment returns a reward and the next state.⁸³ Generally, the goal for the agent is to interact with its environment according to an optimal policy.⁸⁴

The second element of the reinforcement learning framework is the policy. A policy is the way in which an agent makes decisions or chooses actions within a state.⁸⁵ In other words, the agent chooses which action to take when presented with a state based upon the agent's policy.⁸⁶ For example, a greedy person has a policy that routinely guides their decision making toward acquiring the most wealth. The goal of the policy is to allow the agent to advance through the environment so as to maximize a reward.⁸⁷

The third element of the reinforcement learning framework is the reward. Ultimately, the purpose of reinforcement learning is to maximize an agent's reward.⁸⁸ However, the reward itself may be defined by the designer of the algorithm. For each action the agent takes in the environment, a reward is returned.⁸⁹ There are various ways of defining reward based upon the specific application.⁹⁰ But generally, the reward is associated with the final goal of the agent.⁹¹ For example, in a trading algorithm, the reward is money.⁹² In sum, the goal of reinforcement

⁸¹ ALPAYDIN, *supra* note 25, at 126-127.

⁸² SUTTON, BARTO, *supra* note 75, at 2.

⁸³ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 77 (2015).

⁸⁴ *Id.* at 79.

⁸⁵ *Id.*

⁸⁶ SUTTON, BARTO, *supra* note 75, at 39.

⁸⁷ WERBOS, *supra* note 63, at 311.

⁸⁸ SUTTON, BARTO, *supra* note 75, at 50.

⁸⁹ KOCHENDERFER, *supra* note 83, at 77.

⁹⁰ BOSTROM, *supra* note 11, at 239.

⁹¹ MAXIM LAPAN, DEEP REINFORCEMENT LEARNING HANDS-ON 3 (2018).

⁹² *Id.* at 217.

learning is to learn good policies for sequential decision problems, by optimizing a cumulative future reward.⁹³ Interestingly, many thinkers throughout history have argued the human mind is a reinforcement learning system.⁹⁴ And, reinforcement learning algorithms add substantial improvements to deep learning models, especially when the two models are combined.⁹⁵

C. DEEP REINFORCEMENT LEARNING

Deep Reinforcement Learning is an intelligence technique combining deep learning and reinforcement learning. Max Tegmark suggests deep reinforcement learning was developed by Google in 2015.⁹⁶ However, earlier scholarship explores and explains the integration of neural networks in the reinforcement learning paradigm.⁹⁷ Arguably, deep reinforcement learning is a method of general intelligence because of its theoretic capability to solve any continuous control task.⁹⁸ For example, deep reinforcement learning algorithms drive state-of-the-art autonomous vehicles.⁹⁹ However, deep reinforcement learning algorithms show poorer performance on other types of tasks like writing, because mastery of human language is – for now – not describable as a continuous control problem.¹⁰⁰ Regardless of its scalable nature toward general intelligence, deep reinforcement learning is a powerful type of artificial intelligence.¹⁰¹ Generally, there are three different frameworks for deep reinforcement learning: action-value, policy gradient, and actor-critic.¹⁰²

⁹³ Hado van Hasselt, Arthur Guez, and David Silver, *Deep Reinforcement Learning with Q-Learning*, Google DeepMind, 2094 (2018) <https://arxiv.org/abs/1509.06461>.

⁹⁴ WERBOS, *supra* note 63, at 307.

⁹⁵ ALPAYDIN, *supra* note 25, at 136.

⁹⁶ TEGMARK, *supra* note 18, at 85.

⁹⁷ WERBOS, *supra* note 63, at 307.

⁹⁸ M TEGMARK, *supra* note 18, at 39.

⁹⁹ Alex Kendall, et. al., *Learning to Drive in A Day* (2018), <https://arxiv.org/abs/1807.00412>.

¹⁰⁰ NOAM CHOMSKY, SYNTACTIC STRUCTURES 17 (1957).

¹⁰¹ TEGMARK, *supra* note 18, at 39.

¹⁰² Shixun You, et al., *Deep Reinforcement Learning for Target Searching in Cognitive Electronic Warfare*, IEEE Access Vol. 7, 37432, 37438 (2019).

An example of an action-value based framework deep reinforcement learning algorithm is the Deep Q-Network (DQN).¹⁰³ The DQN algorithm is a type of model-free-learning.¹⁰⁴ In model-free-learning, there isn't a formal description of the agent-environment relationship.¹⁰⁵ Instead, the agent randomly explores the environment, gathering information about the environment's states, actions, and rewards.¹⁰⁶ The algorithm stores the information in memory, called experience.¹⁰⁷

The DQN algorithm develops an optimal policy π^* for an agent with a Q-learning algorithm.¹⁰⁸ The optimal policy is the best method of decision making for an agent with the goal of maximizing reward.¹⁰⁹ The Q-learning algorithm maximizes a Q-function: $Q(s, a)$, where s is the state of an environment and a is an action in the state.¹¹⁰ In essence, by applying the optimal Q-function Q^* to every state-action pair (s, a) in an environment, the agent is acting according to the optimal policy.¹¹¹ However, computing $Q(s, a)$ for each state-action pair in environment is computationally expensive.¹¹²

Instead, the DQN algorithm approximates the value of each state state-action pair:¹¹³

$$Q(s, a; \phi) \approx (s, a).$$

¹⁰³ Mnih et al., *supra* note 73, at 529.

¹⁰⁴ KOCHENDERFER, *supra* note 83, at 122.

¹⁰⁵ *Id.* at 121.

¹⁰⁶ LAPAN, *supra* note 91, at 127.

¹⁰⁷ CHARNIAK, *supra* note 56, at 133.

¹⁰⁸ Mnih et al., *supra* note 73, at 529.

¹⁰⁹ KOCHENDERFER, *supra* note 83, at 81.

¹¹⁰ Volodymyr Mnih, Koray Kavukcuoglu, Methods and Apparatus for Reinforcement Learning, U.S. Patent Application No. 14/097,862 at 5 (filed Dec. 5, 2013), <https://patents.google.com/patent/US20150100530A1/en>.

¹¹¹ LAPAN, *supra* note 91, at 144.

¹¹² Volodymyr Mnih, Koray Kavukcuoglu, Methods and Apparatus for Reinforcement Learning, U.S. Patent Application No. 14/097,862 at 5 (filed Dec. 5, 2013), <https://patents.google.com/patent/US20150100530A1/en>.

¹¹³ *Id.*

Here, ϕ represents the function parameters, which are a function's variables.¹¹⁴ The parameters are determined by a neural network using experience replay.¹¹⁵ Experience replay refers to the agent's experiences stored in memory, which are used to train the neural network to approximate the value of state-action pairs.¹¹⁶ The neural network iterates until the convergence of the Q-function determined by the Bellman Equation:¹¹⁷

$$Q^*(s, a) = \mathbb{E}_{s' \sim \varepsilon} \left[r + \gamma \max_{a'} Q^*(s', a') | s, a \right].$$

Here, $\mathbb{E}_{s' \sim \varepsilon}$ refers to the expectation for all states, r is the reward, γ is a discount factor typically defined $0 < \gamma < 1$, allowing present rewards to have higher value.¹¹⁸ Additionally, the *max* function describes an action at which the Q-function takes its maximal value for each state-action pair.¹¹⁹ In other words, the Bellman Equation does two things; it defines the optimal Q-function and allows the agent to consider the reward from its present state as greater relative to similar rewards in future states.¹²⁰

Thus, the DQN algorithm combines Q-learning with a neural network to maximize reward.¹²¹ After the optimal policy is defined according to:

$$\pi^* = Q^*(s', a'),$$

the agent engages in the exploitation of its environment.¹²² During the exploitation phase, the agent maximizes its reward by making decisions according to the optimal policy.¹²³ The DQN is

¹¹⁴ *Id.*

¹¹⁵ CHARNIAK, *supra* note 56, at 133.

¹¹⁶ *Id.*

¹¹⁷ Haney, *supra* note 22 at 162.

¹¹⁸ KOCHENDERFER, *supra* note 83, at 78.

¹¹⁹ Brian S. Haney, *The Optimal Agent: The Future of Autonomous Vehicles & Liability Theory*, 29 ALB. L.J. SCI. & TECH. (2019) (Forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3261275.

¹²⁰ LAPAN, *supra* note 91, at 102-103.

¹²¹ WERBOS, *supra* note 63, at 306-307.

¹²² LAPAN, *supra* note 91, at 127.

¹²³ *Id.*

an off-policy algorithm, meaning it uses data to optimize performance.¹²⁴ Indeed, DQN is essentially a reinforcement learning algorithm, where the agent uses a neural network to decide which actions to take.

A second variant of deep reinforcement learning is the Proximal Policy Optimization (“PPO”) algorithm, a gradient technique.¹²⁵ Similar to the DQN algorithm, the PPO algorithm is a method of model-free learning.¹²⁶ In contrast to the DQN algorithm, PPO is an on-policy algorithm, meaning it does not learn from old data and instead directly optimize policy performance.¹²⁷ One advantage of the PPO model is that it can be used for environments with either discrete or continuous action spaces.¹²⁸

In general, PPO works by computing an estimator of the policy gradient and iterating with a stochastic gradient optimization algorithm.¹²⁹ In other words, the algorithm continuously updates the agent’s policy based on the old policy’s performance.¹³⁰ The PPO update algorithm may be defined:¹³¹

$$\theta_{k+1} = \arg \max_{\theta} \mathbb{E}_{s,a \sim \pi_{\theta_k}} [L(s, a, \theta_k, \theta)].$$

Here, $L(s, a, \theta_k, \theta)$ is the objective function, θ are the policy parameters, θ_k are the policy parameters for k experiment.¹³² Generally, the PPO update is a method of incremental

¹²⁴ Hado van Hasselt, Arthur Guez, and David Silver, *Deep Reinforcement Learning with Q-Learning*, Google DeepMind, 2098 (2018) <https://arxiv.org/abs/1509.06461>.

¹²⁵ John Schulman, et al., *High-Dimensional Continuous Control Using Generalized Advantage Estimation* (2016), <https://arxiv.org/abs/1506.02438>.

¹²⁶ CHARNIAK, *supra* note 56, at 124.

¹²⁷ OpenAI, *Proximal Policy Optimization*, OpenAI Spinning UP (2018) <https://spinningup.openai.com/en/latest/algorithms/ppo.html>.

¹²⁸ *Id.*

¹²⁹ John Schulman, et al., *Proximal Policy Optimization Algorithms*, OpenAI at 2 (2017), <https://arxiv.org/abs/1707.06347>.

¹³⁰ KOCHENDERFER, *supra* note 83, at 80.

¹³¹ OpenAI, *Proximal Policy Optimization*, *supra* note 127.

¹³² *Id.*

improvement for a policy's expected return.¹³³ Essentially, the algorithm takes multiple steps via SGD to maximize the objective.¹³⁴

The PPO algorithm's key to the success is obtaining good estimates of an advantage function.¹³⁵ The advantage function describes the advantage of a particular policy relative to another policy.¹³⁶ For example, if the advantage for the state-action pair is positive, the objective reduces to:¹³⁷

$$L(s, a, \theta_k, \theta) = \min\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)}, (1 + \epsilon)\right) A^{\pi_{\theta_k}}(s, a).$$

Here, $A^{\pi_{\theta_k}}$ is the advantage estimate for the policy given parameters $\pi_\theta(a|s)$, and the hyperparameter ϵ corresponds to how far away the new policy can step from the old while still profiting the objective.¹³⁸ Where the advantage is positive the objective increases and the *min* function puts a limit to how much the objective can increase.¹³⁹

The limitation on the objective increase is called clipping.¹⁴⁰ The algorithm's goal is to make the largest possible improvement on a policy, without stepping so far as to cause performance collapse.¹⁴¹ To achieve this goal, PPO relies on clipping the objective function to remove incentives for the new policy to step far from the old policy.¹⁴² In essence, the clipping serves as a regularizer, minimizing incentives for the policy to change dramatically.¹⁴³

¹³³ Schulman, et al., *supra* note 129, at 2.

¹³⁴ LAPAN, *supra* note 91, at 427.

¹³⁵ Schulman, et al., *supra* note 125.

¹³⁶ OpenAI, Proximal Policy Optimization, *supra* note 127.

¹³⁷ *Id.*

¹³⁸ LAPAN, *supra* note 91, at 432.

¹³⁹ OpenAI, Proximal Policy Optimization, *supra* note 127.

¹⁴⁰ Schulman, et al., *supra* note 129, at 3.

¹⁴¹ OpenAI, Proximal Policy Optimization, *supra* note 127.

¹⁴² Schulman, et al., *supra* note 129, at 3.

¹⁴³ OpenAI, Proximal Policy Optimization, *supra* note 127.

A third variant of Deep Reinforcement Learning and an example of the actor-critic framework is the Deep Deterministic Policy Gradient (“DDPG”) algorithm.¹⁴⁴ Like both DQN and PPO, DDPG is a model-free learning method.¹⁴⁵ However, unlike PPO, DDPG is only applicable in continuous action spaces.¹⁴⁶ In form DDPG is relatively similar to DQN.¹⁴⁷ DDPG is an off-policy algorithm, meaning it re-uses old data.¹⁴⁸ In short, DDPG is a method of deep reinforcement learning using two function approximators, an actor and a critic.¹⁴⁹

The critic estimates the optimal action-value function $a^*(s)$.¹⁵⁰ Generally, the action-value function is tailored to continuous action spaces, defined:

$$a^*(s) = \arg \max_a Q^*(s, a).$$

Here, the optimal action $a^*(s)$ is defined as a value of $Q^*(s, a)$ at which a takes it’s optimal value according to the Bellman Equation.¹⁵¹ The critic’s role is to minimize loss, typically using a means squared error function, or target network, which gives consistent target values.¹⁵² The input of the target network is derived from a replay buffer, utilizing experience replay similar to the DQN algorithm.¹⁵³ As the process occurs, the actor is iteratively updated accordingly.¹⁵⁴ To

¹⁴⁴ LAPAN, *supra* note 91, at 410.

¹⁴⁵ Timothy P. Lillicrap, et al., Continuous Control with Deep Reinforcement Learning, 1 (2016), <https://arxiv.org/abs/1509.02971>.

¹⁴⁶ OpenAI, Deep Deterministic Policy Gradient, OpenAI Spinning UP (2018), <https://spinningup.openai.com/en/latest/algorithms/ddpg.html>.

¹⁴⁷ Lillicrap, et al., *supra* note 145, at 1.

¹⁴⁸ Aprit Agarwal, Katharina Muelling, Katerina Fragkiadaki, Model Learning for Look-ahead Exploration in Continuous Control, Carnegie Mellon University, AAAI (2019), <https://arxiv.org/abs/1811.08086>.

¹⁴⁹ Alex Kendall, et. al., Learning to Drive in A Day (2018), <https://arxiv.org/abs/1807.00412>.

¹⁵⁰ OpenAI, Deep Deterministic Policy Gradient, *supra* note 146.

¹⁵¹ *Id.*

¹⁵² Lillicrap, et al., *supra* note 145, at 2.

¹⁵³ CHARNIAK, *supra* note 56, at 133.

¹⁵⁴ David Silver, et. al. Deterministic Policy Gradient Algorithms, Deem Mind Technologies (2014), <http://proceedings.mlr.press/v32/silver14.pdf>.

learn the optimal policy the DDPG learns a deterministic policy $\pi_\theta(s)$ which gives the action maximizing $Q_\phi(s, a)$:¹⁵⁵

$$\max_{\theta} \mathbb{E}_{s \sim \mathcal{D}} [Q_\phi(s, \pi_\theta(s))].$$

Here, the Q-function parameters Q_ϕ are constants and $s \sim \mathcal{D}$ is the state sampled from the replay buffer.¹⁵⁶

Ultimately, the actor decides which action to take.¹⁵⁷ But, to optimize an agent's reward, after each action, the critic tells the actor defines necessary adjustment for performance improvement.¹⁵⁸ The DDPG algorithm shows promise in continuous control tasks, for robotics control systems.¹⁵⁹ For example, DDPG has shown state-of-the-art success for driving cars.¹⁶⁰ However, the off-policy nature of the algorithm makes it much slower because it takes more computational power to train compared to the PPO and other on-policy algorithms. As computational hardware develops, quantum computers provide a faster method of computing than classical methods.¹⁶¹

In sum, deep learning, reinforcement learning, and deep reinforcement learning provide a framework for analyzing the state-of-the-art in AI technology. While the mathematical models underlying these systems are not new, their capabilities have shown rapid improvement symbiotically with the massive amount of information humans began collecting at the dawn of the digital age.¹⁶² Most importantly, modern AI systems are capable of generalizing about

¹⁵⁵ OpenAI, Deep Deterministic Policy Gradient, *supra* note 146.

¹⁵⁶ OpenAI, Deep Deterministic Policy Gradient, *supra* note 146.

¹⁵⁷ CHARNIAK, *supra* note 56, at 130.

¹⁵⁸ Aleksandra Faust, Oscar Ramirez, et al., PRM-RL: Long-range Robotic Navigation Tasks by Combining Reinforcement Learning and Sampling-based Planning (2018) <https://arxiv.org/abs/1710.03937v2>.

¹⁵⁹ *Id.*

¹⁶⁰ Alex Kendall, et. al., Learning to Drive in A Day (2018), <https://arxiv.org/abs/1807.00412>.

¹⁶¹ Jacob Biamonte, et. al. Quantum Machine Learning 2 (2018) <https://arxiv.org/abs/1611.09347>.

¹⁶² GILDER, *supra* note 79, at 75. *See also* SUSSKIND, *supra* note 31, at 11.

information to make predictions and achieve goals.¹⁶³ As a result, these systems are transforming the foundations of the defense industry, national security threats, and global warfare.

II. SECURITY THREATS

United States National Defense Strategy prioritizes competition with China and Russia.¹⁶⁴ Currently, among these three countries, as well as private corporations, there is an ongoing arms race toward developing the most powerful AI systems.¹⁶⁵ Some hope this continued buildup may be avoided.¹⁶⁶ However, the incentives associated with becoming the world leader in AI technology are great, while the harms of nations falling behind could surely be fatal.¹⁶⁷ Thus, the AI arms races will certainly continue.

Northwestern Law Professor, John McGinnis argues, “The way to think about the effects of AI on war is to think of the consequences of substituting technologically advanced robots for humans on the battlefield.”¹⁶⁸ However, this mode of thought completely fails to communicate AI security threats. Indeed, today the battlefield is everywhere, and the United States is bombarded with cyber-attacks every day.¹⁶⁹ McGinnis further argues “The existential dread of machines that become uncontrollable by humans and the political anxiety about machines’ destructive power on a revolutionized battlefield are overblown.”¹⁷⁰ Yet, China has developed and made publicly available state-of-the-art AI guided missile technology and computer programs.¹⁷¹ And, Russia routinely and intentionally manipulates United States voters on social

¹⁶³ TEGMARK, *supra* note 18, at 85-86.

¹⁶⁴ Mark D. Miles and Charles R. Miller, Global Risks and Opportunities the Great Power Competition Paradigm, JFQ 94, 3rd Quarter, at 80 (2019).

¹⁶⁵ An Interview with Paul M. Nakasone, National Defense University Press, JFQ 92, 1st Quarter, at 5 (2019).

¹⁶⁶ Baker, *supra* note 12, at 5.

¹⁶⁷ BOSTROM, *supra* note 11, at 96-97.

¹⁶⁸ John O. McGinnis, *Accelerating AI*, 104 NW. U. L. REV. 1253, 1266 (2010).

¹⁶⁹ John P. Carlin, *Detect, Disrupt, Deter: A Whole-of-Government Approach to National Security Cyber Threats*, 7 HARV. NAT’L SEC. J. 391, 398 (2016).

¹⁷⁰ McGinnis, *supra* note 168, at 1254.

¹⁷¹ Shixun You, et al., *supra* note 102, at 37447.

media for the purposes of influencing political elections.¹⁷² In short, AI is the most important weapon in modern warfare, defense, and national security. This Part will discuss AI applications in three types of security threats: missile attack, cyber-attack, and general intelligence.

A. MISSILES

Richmond School of Law Professor and Member of the Center for New American Security's Task Force on Artificial Intelligence and National Security, Rebecca Crootof, suggests weapons may be grouped into three categories: inert, semi-autonomous, and autonomous.¹⁷³ Inert weapons require human operation to be lethal, including stones, knives, or handheld firearms.¹⁷⁴ Semi-autonomous weapon systems, have autonomous capabilities in functions relevant to target selection and engagement, but cannot both select and engage targets independently.¹⁷⁵ Third, autonomous weapon systems are capable of independently selecting and engaging targets based on conclusions derived from gathered information and preprogramed constraints.¹⁷⁶

Professor Crootof argues “autonomous weapon systems in use today act in largely predictable ways.”¹⁷⁷ Similarly, The Honorable James E. Baker, argues autonomous weapons systems are nothing new.¹⁷⁸ Judge Baker claims autonomous weapons have been standard military technology since the 1970s, and the United States reserves the technology for defensive purposes.¹⁷⁹ Further, according to the Department of Defense, American “adversaries are also developing an increasingly diverse, expansive, and modern range of offensive missile systems

¹⁷² U.S. Department of Justice, Report on The Investigation into Russian Interference in the 2016 Presidential Election, Vol I at 4 (March 2019), https://www.justice.gov/storage/report_volume1.pdf.

¹⁷³ Rebecca Crootof, Autonomous Weapons Systems and the Limits of Analogy, 9 Harv. Nat'l Sec. J. 51, 59 (2018).

¹⁷⁴ *Id.*

¹⁷⁵ *Id.*

¹⁷⁶ *Id.*

¹⁷⁷ *Id.* at 60.

¹⁷⁸ Baker, *supra* note 12, at 3.

¹⁷⁹ *Id.*

that can threaten U.S. forces abroad.”¹⁸⁰ However, these perspectives sincerely underestimate the capabilities modern missile systems, particularly in light of AI advancements.¹⁸¹ Inarguably, AI has changed the role of robotics control systems in warfare.¹⁸²

It is important to understand foreign adversaries have the ability to attack the United States homeland with AI controlled missile systems at such a scale to which the United States would be entirely unable to respond.¹⁸³ Indeed, in a recent study funded by the National Natural Science Foundation of China, *Deep Reinforcement Learning for Target Searching in Cognitive Electronic Warfare* (China AI Missile Study), researchers demonstrate Chinese capabilities in deep reinforcement learning control systems for missile control.¹⁸⁴ The United States funded similar research through the Naval Post-Graduate School in a 2017 report, *A Framework Using Machine Vision and Deep Reinforcement Learning for Self-Learning Moving Objects in a Virtual Environment* (Navy AI Study).¹⁸⁵ However, the Chinese research is not only far more advanced, but also open-sourced.¹⁸⁶ Indeed, China’s system is adaptable to any environment or target across the globe.¹⁸⁷ And, the code for China’s deep reinforcement learning missile control systems is available on GitHub.¹⁸⁸ Further, Google’s TensorFlow, is also available open-source and designed specifically for manufacturing and scalability.¹⁸⁹

¹⁸⁰ Department of Defense, Missile Defense Review, Executive Summary III (2009), https://www.defense.gov/Portals/1/Interactive/2018/11-2019-Missile-Defense-Review/The%202019%20MDR_Executive%20Summary.pdf.

¹⁸¹ Shixun You, et al., *supra* note 102, at 37447.

¹⁸² JOHN JORDAN, ROBOTS 133 (2016).

¹⁸³ Shixun You, et al., *supra* note 102, at 37447.

¹⁸⁴ *Id.* at 37434.

¹⁸⁵ Richard Wu, et al., *supra* note 37.

¹⁸⁶ Shixun You, et al., *supra* note 102, at 37435.

¹⁸⁷ *Id.* at 37441.

¹⁸⁸ youshixun, vCEW New model of cognitive electronic warfare with countermeasures, GitHub <https://github.com/youshixun/vCEW> (2019).

¹⁸⁹ TensorFlow 2.0 Alpha is Available, TensorFlow, (2019), <https://www.tensorflow.org/install>.

AI missile technology is comparatively simple relative to AI controlled vehicles or rocket boosters due to the general lack of obstacles in a missile's environment. Indeed, there are at most three elements needed to control an AI missile. First, a means of perception, which is commonly achieved with Light Detection and Ranging Device (LiDAR) sensors.¹⁹⁰ LiDAR sensors simply work by sending light pulses from a transmitter and measuring return times with a receiver.¹⁹¹ The time it takes for a pulse to return measures distance according to $\frac{tc}{2} = d$, where t is travel time, c is the speed of light, and d is the distance between the LIDAR sensor and the object.¹⁹² The receiver then generates a point cloud map of the environment for processing.¹⁹³

Second, the processing typically occurs with convolutional neural networks (CNNs), which show state-of-the-art performance in computer vision tasks.¹⁹⁴ CNNs utilize convolutional mathematics to perform computer vision tasks like object detection and classification.¹⁹⁵ Further, CNNs are well suited for three-dimensional point cloud environments and integration with reinforcement learning algorithms.¹⁹⁶ One study, conducted by research firm OpenAI, demonstrated the effectiveness of CNNs in real-time obstacle detection when integrated with reinforcement learning systems.¹⁹⁷

The third element is a method of optimization for decision making, commonly reinforcement learning.¹⁹⁸ For example, the China AI Missile Study explored the use of DQN,

¹⁹⁰ Jeff Hect, *Lidar for Self-Driving Cars*, Optics & Photonics News, 29 (January 2018), https://www.osa-opn.org/home/articles/volume_29/january_2018/features/lidar_for_self-driving_cars/.

¹⁹¹ Gaetan Pennecot, et. al., *Devices and Methods for a Rotating LIDAR Platform with Shared Transmit/Receive Path*, Google, Inc., No: 13/971,606, (Aug. 20, 2013), <https://patents.google.com/patent/US9285464B2/en>.

¹⁹² Matthew J. McGill, *LIDAR Remote Sensing*, NASA Technical Reports Server (NTRS) (2002).

¹⁹³ Maturana, Scherer, *supra* note 58, at 2.

¹⁹⁴ *Id.*

¹⁹⁵ Legrand, *supra* note 55, at 23.

¹⁹⁶ Mnih et al., *supra* note 73, at 530.

¹⁹⁷ Gregory Kahn, *Uncertainty-Aware Reinforcement Learning for Collision Avoidance* (2017).

¹⁹⁸ Shixun You, *Completing Explorer Games with a Deep Reinforcement Learning Framework Based on Behavior Angle Navigation*, Electronics at 17 (2019) <https://www.mdpi.com/2079-9292/8/5/576>.

PPO, and DDPG for control in its simulated, real-time physics engine.¹⁹⁹ And, the Navy AI Missile Study experimented with the DQN algorithm.²⁰⁰ In the context of missile control, the reinforcement learning agent is able to visualize its environment with LiDAR and a CNN and generalize to avoid obstacles, including defense missiles.²⁰¹ This framework maximizes the probability of success in target searching, detection, and engagement regardless of motion dynamics.²⁰² As such, AI missile systems guided by LiDAR sensor data and controlled with deep reinforcement learning algorithms have the capability to attack any target on Earth, or in the atmosphere, with pixel precision.²⁰³ Importantly, this information and the tools to build such a system are widely available on the internet.²⁰⁴

In short, Professor Crootof and Judge Baker's misunderstandings about the nature of autonomous weapons derive from their grouping of all autonomous weapons as having analogous abilities and posing analogous levels of threat.²⁰⁵ Indeed, modern AI missile systems, specifically, deep reinforcement learning systems do not act in the same predictable ways as the autonomous missile systems of the 1970s.²⁰⁶ In fact, they are much different than the autonomous weapons of the 1970s.²⁰⁷ Critically, deep reinforcement learning missiles today are able to generalize about their environment, adapting, and evolving with the battlefield.²⁰⁸

¹⁹⁹ Shixun You, et al., *supra* note 102, at 37438.

²⁰⁰ Richard Wu, et al., *supra* note 37, at 233.

²⁰¹ Gregory Kahn, *Uncertainty-Aware Reinforcement Learning for Collision Avoidance* (2017).

²⁰² Serena Yung, et al., Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos (2017) <https://arxiv.org/abs/1507.05738>.

²⁰³ Shixun You, et al., *supra* note 102, at 37438.

²⁰⁴ See Richard Wu, et al., *supra* note 37. See also youshixun, vCEW New model of cognitive electronic warfare with countermeasures, GitHub, <https://github.com/youshixun/vCEW> (2019). See also TensorFlow 2.0 Alpha is Available, TensorFlow, (2019), <https://www.tensorflow.org/install>. See also Shixun You, et al., *supra* note 102, at 37434-37447.

²⁰⁵ Crootof, *supra* note 172 at 59.

²⁰⁶ *Id.* at 60.

²⁰⁷ Baker, *supra* note 12, at 3. See also Shixun You, Completing Explorer Games with a Deep Reinforcement Learning Framework Based on Behavior Angle Navigation, *Electronics* at 17 (2019) <https://www.mdpi.com/2079-9292/8/5/576>.

²⁰⁸ Richard Wu, et al., *supra* note 37, at 231.

Specifically, Chinese AI missile technology “is enhanced by the powerful generalization ability of ... deep convolutional neural network[s].”²⁰⁹

Indeed, according to the 2019 Department of Defense Missile Review, China now has the ability to threaten the United States with about 125 nuclear missiles.²¹⁰ The Review explains while the United States relies on deterrence to protect against sophisticated actors like China, active U.S. missile defense efforts must outpace rogue offensive missile strikes.²¹¹ But, despite massive over-spending, the United States Military is falling behind.²¹² Further, due to the availability of AI missile system construction and control information across the internet, AI missile attacks are an immediate national security threat.²¹³ This is especially true considering hostile relationships with Iran, North Korea, and Middle Eastern militant organizations.²¹⁴ Intimately related to the applications of AI missiles in national security, the role of AI in cybersecurity attacks is also a cause for immediate concern.

B. CYBERSECURITY

Top U.S. national security officials believe that AI and machine learning will have transformative implications for cybersecurity and cyber war.²¹⁵ For example, according to National Security Agency (NSA) Director Paul M. Nakson, “As we look at near-peer competitors, China and Russia clearly are at the top of the list because they have the capacity to

²⁰⁹ Shixun You, et al., *supra* note 102, at 37438.

²¹⁰ Department of Defense, Missile Defense Review, *supra* note 180.

²¹¹ *Id.*

²¹² Craig Whitlock, Bob Woodward, Pentagon buries evidence of \$125 billion in bureaucratic waste, The Washington Post (December 5, 2016) https://www.washingtonpost.com/investigations/pentagon-buries-evidence-of-125-billion-in-bureaucratic-waste/2016/12/05/e0668c76-9af6-11e6-a0ed-ab0774c1eaa5_story.html. See also NATO, Defence Expenditure of NATO Countries (2012-2019) (June 25, 2019), https://www.nato.int/nato_static_fl2014/assets/pdf/pdf_2019_06/20190625_PR2019-069-EN.pdf.


²¹³ See youshixun, vCEW New model of cognitive electronic warfare with countermeasures, GitHub <https://github.com/youshixun/vCEW> (2019). See also TensorFlow 2.0 Alpha is Available, TensorFlow, (2019), <https://www.tensorflow.org/install>.

²¹⁴ Department of Defense, Missile Defense Review, *supra* note 180.

²¹⁵ Greg Allen, Taniel Chan, *Artificial Intelligence and National Security*, Harvard Kennedy School Belfer Center for Science and International Affairs, 18 (July 2017).

operate across the full spectrum of cyberspace operations.”²¹⁶ However, Nakasone seems to be behind the times in terms of AI military applications in cyberspace. Indeed, in a 2019 interview with National Defense University, Nakasone stated, “I think that the early instantiation of AI will be on the defensive side.”²¹⁷ Further, recent cybersecurity data and research provide strong evidence AI is already fueling offensive state sponsored cyber-attacks.

The Center for Strategic & International Studies (CSIS), a policy research organization, records reported significant cyber-attacks on governments, defense and high technology companies, or economic crimes in excess of a million dollars.²¹⁸ The table below illustrates the number of Chinese and Russian cyber-attacks per year.²¹⁹



²¹⁶ An Interview with Paul M. Nakasone, National Defense University Press, JFQ 92, 1st Quarter, at 5 (2019).

²¹⁷ *Id.* at 8.

²¹⁸ Center for Strategic and International Studies, Significant Cyber Incidents (August 2019), <https://www.csis.org/programs/technology-policy-program/significant-cyber-incidents>.

²¹⁹ See Appendix B. Significant Cyber Incidents for chart data (Chart prepared by author with information from, Center for Strategic and International Studies, Significant Cyber Incidents (August 2019) <https://www.csis.org/programs/technology-policy-program/significant-cyber-incidents>).

While the accelerating trend in cyber-attacks could be due to a number of factors, AI is unquestionably playing a contributing role. Indeed, AI provides more powerful attack mechanisms which are widely available, making this trend likely to continue.²²⁰

Military defense and political hacking are now commonplace.²²¹ Some of the most significant defense and political related attacks reported by the CSIS include:²²²

1. *May 2019*. Hackers affiliated with the Chinese intelligence service reportedly had been using NSA hacking tools since 2016, more than a year before those tools were publicly leaked.
2. *March 2019*. U.S. officials reported that at least 27 universities in the U.S. had been targeted by Chinese hackers as part of a campaign to steal research on naval technologies.
3. *October 2018*. Media reports state that U.S. agencies warned President Trump that that China and Russia eavesdropped on call made from an unsecured phone.
4. *August 2018*. Microsoft announces that Russian hackers had targeted U.S. Senators and Center conservative think tanks critical of Russia.
5. *February 2016*. Hackers breached the U.S. Department of Justice's database, stealing and releasing the names, phone numbers, and email addresses of 30,000 DHS and FBI employees.
6. *October 2016*. The U.S. Director of National Intelligence and Department of Homeland Security jointly identified Russia as responsible for hacking the Democratic National Committee and using WikiLeaks to dump emails obtained in the hack.
7. *April 2015*. U.S. officials report that hackers gained access to White House networks and sensitive information, such as "real-time non-public details of the president's schedule," through the State Department's network, which has had continued trouble in ousting attackers.
8. *February 2012*. Media reports say that Chinese hackers stole classified information about the technologies onboard F-35 Joint Strike Fighters.²²³

The pervasive use of cyberspace for a variety of covert international operations provides support for the argument that an AI arms race is in full swing.²²⁴ Indeed, AI provides a decisive

²²⁰ Hyrum S. Anderson, et.al., *Learning to Evade Static PE Machine Learning Malware Models via Reinforcement Learning*, Cornell University Library (2018), <https://arxiv.org/abs/1801.08917>. See also endgameinc, *Malware Env for OpenAI Gym* (2018), <https://github.com/endgameinc/gym-malware>.

²²¹ Center for Strategic and International Studies, *Significant Cyber Incidents* (August 2019), <https://www.csis.org/programs/technology-policy-program/significant-cyber-incidents>.

²²² *Id.*

²²³ *Id.* (List derived from preceding source).

²²⁴ Ido Kilovaty, *Doxfare: Politically Motivated Leaks and the Future of the Norm on Non-Intervention in the Era of Weaponized Information*, 9 *Harv. Nat'l Sec. J.* 146, 150 (2018).

advantage in the penetration and security of networks.²²⁵ While it is difficult to know the specific architectures used, the wide-spread availability of AI cyber-attack research and code strongly suggests malicious AI software is already common-place.²²⁶

As the data reflects, cyber-dependent nations are vulnerable to political manipulation.²²⁷ Indeed, politically motivated cyber-attacks are far from a new phenomenon.²²⁸ And, Russia has significantly expanded its budget in these areas over the last few years to sway public opinion around the world.²²⁹ One direct example, includes the Internet Research Agency (IRA), a Russian intelligence company, and The Main Intelligence Directorate of the Russian Army's (GRU) hacking operations during the 2016 Presidential Election.²³⁰ According to the Mueller Report, the IRA and GRU both engaged in hacking Democratic National Committee networks for the purpose of influencing the election's outcome.²³¹ Further, the *Mueller Report* details the IRA's social media tactics with the intent of swaying voters.²³² Another example, includes the Cambridge Analytica Scandal, where 87 million Facebook users had their personal data exposed without their consent and used by Cambridge Analytica to support political campaigns.²³³ As a result, data driven AI cyber-tools for manipulating voters via social media are critical to the integrity of United States elections and thus, National Security.²³⁴

²²⁵ Alberto Perez Veiga, *Applications of Artificial Intelligence to Network Security*, 14 (2018), <https://arxiv.org/abs/1803.09992>.

²²⁶ Miles Brundage, et. al. *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*, 34 (Cornell University Library 2018) <https://arxiv.org/pdf/1802.07228.pdf>. See also Anderson, et.al., *supra* note 220.

²²⁷ Kilovaty, *supra* note 224, at 150.

²²⁸ WERBOS, *supra* note 63, at 184.

²²⁹ Kilovaty, *supra* note 224, at 158.

²³⁰ U.S. Department of Justice, *supra* note 171, at 4.

²³¹ *Id.*

²³² *Id.* at 14.

²³³ Felipe González, et al., Global Reactions to the Cambridge Analytica Scandal: An Inter-Language Social Media Study 1 (2019) https://faculty.washington.edu/aragon/pubs/LA_WEB_Paper.pdf.

²³⁴ Jon M. Garon, *Cyber-World War III: Origins*, 7 J.L. & CYBER WARFARE 1, 9 (2018).

Unsurprisingly, AI research in cybersecurity is rapidly expanding. For example, one piece of scholarship specifically details guidelines for the development of malicious AI software.²³⁵ The scholarship was written to demonstrate the practical possibilities of developing malicious machine learning algorithms capable of harming humans.²³⁶ Another paper, *Adversarial Reinforcement Learning in a Cyber Security Simulation*, introduces a game played between adversarial reinforcement learning systems, an attack and a defender.²³⁷ The paper ultimately introduced a cyber security simulation with two Markov agents playing against each other as attacker and defender of a cyber network.²³⁸ Drawing on this work a third paper was introduced in 2019, *Deep Reinforcement Learning for Cyber Security*, which developed a similar simulation, exploring the effectiveness of action-value, policy, and actor-critic deep reinforcement learning models for network attack and defense.²³⁹

Interestingly, one paper, *Learning to Evade Static PE Machine Learning Malware Models via Reinforcement Learning*, specifically detailed reinforcement learning malware models and open-sourced the code on GitHub.²⁴⁰ The reinforcement learning model in the paper is an evasive malware variant, which can be effective on samples not used during training.²⁴¹ Samples, are prior examples of malware detection software for processing and classifying data.²⁴² The purpose of training is to allow the reinforcement learning agent to generalize about

²³⁵ Federico Pistono, Roman Yampolskiy, *Unethical Research: How to Create Malevolent Artificial Intelligence* (2016), <https://arxiv.org/pdf/1605.02817.pdf>.

²³⁶ *Id.*

²³⁷ Richard Elderman, et al., *Adversarial Reinforcement Learning in a Cyber Security Simulation* (2017), <https://pdfs.semanticscholar.org/b495/7266c43f1c76ed8dbb275b8510f6fb1e063c.pdf>.

²³⁸ *Id.* at 566.

²³⁹ Thanh Thi Nguyen and Vijay Janapa Reddi, *Adversarial Reinforcement Learning in a Cyber Security Simulation* 4 (2019), <https://arxiv.org/abs/1906.05799>.

²⁴⁰ Anderson, et.al., *supra* note 220. *See also* endgameinc, *Malware Env for OpenAI Gym* (2018), <https://github.com/endgameinc/gym-malware>.

²⁴¹ Anderson, et.al., *supra* note 220.

²⁴² Andrea Scripa Els, *Artificial Intelligence as A Digital Privacy Protector*, 31 Harv. J.L. & Tech. 217, 221 (2017).

new experiences; in other words avoid new malware detection software.²⁴³ The algorithm uses Q-learning and an MDP framework to train a reinforcement learning agent, where the reward function is associated with an anti-malware detection system.²⁴⁴ If the anti-malware system detects the agent, its reward is zero, otherwise its reward is ten.²⁴⁵ The researchers used OpenAI Gym, a software tool for building reinforcement learning environments, to conduct their experiments.²⁴⁶ In sum, the key contribution of the paper is that the attacker requires no prior knowledge about the target to successfully penetrate a system under attack.²⁴⁷

Cyberspace allows humanity to communicate, trade, research, and share information on a global scale.²⁴⁸ Yet, cyber-attacks are surging and increasing in volume each year.²⁴⁹ As a result, data breaches have escalated increasing criticisms from regulators, private plaintiffs, and public opinion.²⁵⁰ On the global scale, data and information are only as secure as the algorithms and network structures by which they are protected. Arguably, nothing that happens on a computer is secure, let alone private.²⁵¹ Thus, there is a great deal of uncertainty surrounding security at all levels.²⁵² Indeed, the weaponization of cyberspace and resulting cyberwarfare, are creating a world in which traditional security defense models are futile.²⁵³ And, the capability of hacking an enemy car, plane, nuclear reactor, communication system, or financial system has the potential to cripple an opposition's economy and defense capability.²⁵⁴

²⁴³ Anderson, et.al., *supra* note 220.

²⁴⁴ *Id.*

²⁴⁵ *Id.*

²⁴⁶ Greg Brockman, et al., OpenAI Gym (2016), <https://arxiv.org/abs/1606.01540>.

²⁴⁷ Anderson, et.al., *supra* note 220.

²⁴⁸ Kilovaty, *supra* note 224, at 179.

²⁴⁹ Loren F. Selznick, Carolyn LaMacchia, *Cybersecurity Liability: How Technically Savvy Can We Expect Small Business Owners to Be?*, 13 J. Bus. & Tech. L. 217, 219 (2018).

²⁵⁰ Thad A. Davis, et. al., *The Data Security Governance Conundrum: Practical Solutions and Best Practices for The Boardroom and The C-Suite*, 2015 Colum. Bus. L. Rev. 613, 615 (2015).

²⁵¹ Jack M. Balkin, *The Constitution in The National Surveillance State*, 93 MINN. L. REV. 1, 14 (2008).

²⁵² TIM MATHER, ET AL., CLOUD SECURITY AND PRIVACY 31 (2009).

²⁵³ NAYAN B. RUPARELIA, CLOUD COMPUTING 108 (2016).

²⁵⁴ TEGMARK, *supra* note 18, at 118.

From a national security perspective, the United States' is no longer the world leader in cybersecurity.²⁵⁵ Indeed, The United States is under constant attack online.²⁵⁶ Adversaries have caught up and arguably surpassed domestic cyber capabilities.²⁵⁷ Many nations continue investing heavily in AI research and its commercial capabilities, exploiting new advances.²⁵⁸ At the same time, the United States has been ignoring AI over the last decade.²⁵⁹ As a result, the United States is consistently attacked on its vulnerable cyber-front and subject to imminent national security threats.²⁶⁰ While the United States faces imminent cybersecurity threats stemming from AI, the development of Artificial General Intelligence poses threats to both national and global security.

C. ARTIFICIAL GENERAL INTELLIGENCE

AI's holy grail, Artificial General Intelligence (AGI), is a system capable of achieving any goal.²⁶¹ Current methods of AGI development are commonly referred to as whole-brain emulation, where the idea is to reverse engineer the human brain with computation.²⁶² Yet, the Wright brothers' first flight was not aboard a mechanical bird with flapping wings.²⁶³ In fact, prominent physicist Michio Kaku argues computers cannot truly replicate the behavior of human brains.²⁶⁴ Kaku argues the shortcomings of modern neural networks persuasively, focusing on their inability to account from neurochemical fluctuations in information transfer.²⁶⁵

²⁵⁵ Kadtke, Wharton, *supra* note 2, at 1.

²⁵⁶ John P. Carlin, *Detect, Disrupt, Deter: A Whole-of-Government Approach to National Security Cyber Threats*, 7 HARV. NAT'L SEC. J. 391, 398 (2016).

²⁵⁷ An Interview with Paul M. Nakasone, National Defense University Press, JFQ 92, 1st Quarter, at 5 (2019).

²⁵⁸ Kadtke, Wharton, *supra* note 2, at 2.

²⁵⁹ *Id.*

²⁶⁰ Center for Strategic and International Studies, Significant Cyber Incidents (August 2019), <https://www.csis.org/programs/technology-policy-program/significant-cyber-incidents>.

²⁶¹ TEGMARK, *supra* note 18, at 68.

²⁶² KURZWEIL, *supra* note 9, at 124.

²⁶³ TEGMARK, *supra* note 18, at 156.

²⁶⁴ MICHIO KAKU, THE FUTURE OF THE MIND 342 (2014).

²⁶⁵ *Id.*

However, the legendary machine learning developer Paul John Werbos argued, from an engineering point of view, the human brain is an information processing system.²⁶⁶ Therefore, it may be more likely AGI will result from a more simplified neural processing model capable of recursive self-improvement.²⁶⁷ For example, famed philosopher of mind, Zoltan Torey approaches the mind from a linguistic perspective.²⁶⁸ Indeed, according to Torey, the mind is made up of perceptions and words corresponding to those perceptions.²⁶⁹

Yet, some argue that AGI may never happen.²⁷⁰ For example, the late Microsoft co-founder Paul Allen argues that scientific progress is irregular and hypothesizes that at the end the twenty-first century humans will have yet to achieve AGI.²⁷¹ Indeed, current systems are far from achieving many goals, particularly time-consuming tasks.²⁷² One example of such a task would be for a system to litigate a complex case in court from the filing of the complaint, through discovery, all the way to trial and verdict.²⁷³

To date, the closest mankind has come toward developing an AGI was Volodymyr Mnih's seminal paper, *Human-Level Control Through Deep Reinforcement Learning*, where Mnih introduces the DQN algorithm and associated software code for playing Atari Games.²⁷⁴ Max Tegmark remarked of Minh's paper, "deep reinforcement learning is a completely general technique."²⁷⁵ In this sense, Mnih's algorithm, the DQN, generalizes about its environment to

²⁶⁶ WERBOS, *supra* note 63, at 305.

²⁶⁷ MURRAY SHANAHAN, THE TECHNOLOGICAL SINGULARITY 151 (2015). *See also* TEGMARK, *supra* note 18, at 156.

²⁶⁸ ZOLTAN TOREY, THE CONSCIOUS MIND 61 (2014).

²⁶⁹ *Id.*

²⁷⁰ Paul G. Allen, *The Singularity Isn't Near*, MIT Technology Review (2011)

<https://www.technologyreview.com/s/425733/paul-allen-the-singularity-isnt-near/>.

²⁷¹ Paul G. Allen, *The Singularity Isn't Near*, MIT Technology Review (2011)

<https://www.technologyreview.com/s/425733/paul-allen-the-singularity-isnt-near/>.

²⁷² MURRAY SHANAHAN, THE TECHNOLOGICAL SINGULARITY 5 (2015).

²⁷³ This example ignores any legal ethics issues and is simply meant to be illustrative of a complicated task.

²⁷⁴ Mnih et al., *supra* note 73, at 529. *See also* Code for Human-Level Control through Deep Reinforcement Learning (2015), <https://sites.google.com/a/deepmind.com/dqn/>.

²⁷⁵ TEGMARK, *supra* note 18, at 85.

achieve its goal.²⁷⁶ But the DQN is limited by its environment, static reward structure, and training. Thus, a challenge exists to improve the generalizable qualities of current state-of-the-art AI systems.

From a national security perspective AGI is the end-all-be-all in advanced weaponry.²⁷⁷ Any state or corporation capable of controlling AGI would surely be capable of conquering the world.²⁷⁸ Indeed, with control of a system capable of achieving any goal controlling enemy defense systems, manipulating public opinion, and controlling information networks would be relatively simple.²⁷⁹ However, there exists a question as to whether a human creator could control an AGI.²⁸⁰ According to Max Tegmark, “we have no idea what will happen if humanity succeeds in building human-level AGI.”²⁸¹ Thus, we cannot take for granted that the outcome will be positive if AGI is created.²⁸²

III. POLICY

New generations of advanced technologies are changing the power dynamics of our global society.²⁸³ Yet, legal scholarship on the topic of AI policy has denied and relatively ignored the national security threats associated with AI’s weaponization.²⁸⁴ For example, University of Washington Law Professor, Ryan Calo encourages regulators not to be distracted by claims of an “AI Apocalypse” and to focus their efforts on “more immediate harms.”²⁸⁵ However, it is important to realize, AI’s most immediate applications will be in warfare.

²⁷⁶ Mnih et al., *supra* note 73, at 531. *See also* Code for Human-Level Control through Deep Reinforcement Learning (2015), <https://sites.google.com/a/deepmind.com/dqn/>.

²⁷⁷ BOSTROM, *supra* note 11, at 106-107.

²⁷⁸ *Id.* at 96-97.

²⁷⁹ TEGMARK, *supra* note 18, at 118.

²⁸⁰ BOSTROM, *supra* note 11, at 155.

²⁸¹ TEGMARK, *supra* note 18, at 156.

²⁸² PETER THEIL, *ZERO TO ONE* 195 (2014).

²⁸³ Kadtko, Wharton, *supra* note 2, at 2.

²⁸⁴ Calo, *supra* note 14, at 432.

²⁸⁵ *Id.* at 431.

Generally, it is accepted law never keeps up with technology.²⁸⁶ However, the kinetics of the two systems are relative and it's more of an apples to oranges comparison. More likely to be true is that United States policy makers and military leaders are ill-equipped to put policies in place to maintain military superiority. For example, Judge Baker explains, "I do not feel a sense of urgency to address the legal, ethical, and policy challenges ahead."²⁸⁷ Another example includes South Carolina Senator Lindsay Graham's infamous question to Mark Zuckerberg, "Is Twitter the same as what you do?" during the Senate Judiciary & Commerce Committees Joint Hearing on Facebook Data Use.²⁸⁸ As Elon Musk persuasively argues, what governments need right now is not oversight, but insight because right now the Government does not even have insight into AI issues.²⁸⁹ Specifically, Musk contends we need technically capable people in government positions who can monitor AI's progress and steer it if warranted.²⁹⁰ This Part explores the policies and developments from the three countries leading the way in AI militarization, Russia, China, and the United States.

Professor Crootof argues in any armed conflict, the right of the parties to the conflict to choose methods or means of warfare is not unlimited.²⁹¹ And, both customary international law and various treaties circumscribe which weapons may be lawfully fielded.²⁹² However, this line of argument does not apply in the context of AI. In fact, international laws and treaties are not

²⁸⁶ Baker, *supra* note 12, at 7.

²⁸⁷ *Id.*

²⁸⁸ Committee on the Judiciary, Senate Committee on the Judiciary, Senate Committee on Commerce, Science, and Transportation, Facebook, Social Media Privacy, and the Use and Abuse of Data (April 10, 2018), <https://www.judiciary.senate.gov/meetings/facebook-social-media-privacy-and-the-use-and-abuse-of-data>. (Senator Graham Questioning Zuckerberg at 1:53:40-1:53:51).

²⁸⁹ *Elon Musk at the National Governors Association 2017 Summer Meeting*, C-SPAN (July 15, 2017), <https://www.c-span.org/video/?431119-6/elon-musk-addresses-nga>. (Musk responding to Arizona Governor Doug Ducey at 57:00-60:00).

²⁹⁰ *Elon Musk at the National Governors Association 2017 Summer Meeting*, C-SPAN (July 15, 2017), <https://www.c-span.org/video/?431119-6/elon-musk-addresses-nga>. See also TEGMARK, *supra* note 18, at 108.

²⁹¹ Crootof, *supra* note 172 at 62.

²⁹² *Id.*

laws in the sense they are not enforceable because the nature of law rests on the assumption certain conduct be binding.²⁹³ As Hart argued, “If the rules of international law are not binding it is surely indefensible to take seriously their classification as law.”²⁹⁴ The Latin maxim *Auctoritas non veritas facit legem*; which stands for the principle, authority, not truth, makes law, provides insight into the fickle nature of international law.²⁹⁵ Or, in the words the English poet John Lyly, “All is fair in love and war.”²⁹⁶ Therefore, any notion of an international AI treaty would be moot. In addition to the United States, China and Russia, are making significant investments in AI for military purposes.²⁹⁷

A. CHINA

In July 2017, China’s State Council, released an AI plan and strategy calling for China to pass the United States by 2020 and become the world’s leader in AI by 2030, committing \$150 billion to the goal.²⁹⁸ By the end of 2018, Chinese leadership assessed the program’s development as surpassing the United States, achieving its objective earlier than expected.²⁹⁹ A key advantage of China’s recent strategy has been in the development of innovative new systems, in direct contrast to the United States, whose commitments remain to updating outdated technologies and political favors toward the military industrial complex.³⁰⁰

Indeed, as a direct result of China’s recent investments, China’s military and intelligence services possess the sophistication and resources to hack network systems, establish footholds

²⁹³ H.L.A. HART, *THE CONCEPT OF LAW* 214 (3d. 2012).

²⁹⁴ *Id.*

²⁹⁵ Thomas Hobbes: A Pioneer of Modernity, 9 (2015) <https://www.sunypress.edu/pdf/63242.pdf>.

²⁹⁶ John Lyly Quotes, goodreads, (2019) https://www.goodreads.com/author/quotes/139084.John_Lyly.

²⁹⁷ Department of Defense, Summary of the 2018 Department of Defense Artificial Intelligence Strategy 5 (2018), <https://www.defense.gov/Newsroom/News/Article/Article/1755942/dod-unveils-its-artificial-intelligence-strategy/>.

²⁹⁸ Baker, *supra* note 12, at 1.

²⁹⁹ Allen, *supra* note 6, at 9.

³⁰⁰ Alexander Rogosa, *Shifting Spaces: The Success of The SpaceX Lawsuit and The Danger of Single-Source Contracts in America’s Space Program*, 25 Fed. Circuit B.J. 101, 104 (2015). See also Chris Anderson, Elon Musk’s Mission to Mars, *Wired Magazine* (October 21, 2012) <https://www.wired.com/2012/10/ff-elon-musk-qa/>.

behind perimeter defenses, exfiltrate valuable information, and sabotage critical network functions.³⁰¹ In fact, Chinese government organizations routinely translate, disseminate, and analyze U.S. government and think tank reports about AI.³⁰² Further, in 2017, China expressed a desire to utilize AI for flight guidance and target recognition systems in new generations of its own cruise missiles.³⁰³ Just two years later, that desire was realized and China is the world's leader in missile technology, with the development of deep reinforcement learning control systems for targeting and guidance.³⁰⁴ Further, Chinese intercontinental ballistic missile and cruise missile systems reflect the state-of-the-art.³⁰⁵

Chinese commercial markets for autonomous drones and AI surveillance technologies have seen significant growth and success.³⁰⁶ Additionally, Chinese weapons manufacturers already are selling armed AI controlled drones.³⁰⁷ Chinese AI market success directly increases its military and intelligence abilities because Chinese companies developing AI work in close cooperation with the Chinese Military.³⁰⁸ Some argue many Chinese AI achievements are actually achievements of multinational research teams and companies.³⁰⁹ For example, regarding SpaceX's decision not to patent its rocket technologies, Founder & CEO Elon Musk stated, "our primary long-term competition is China – if we published patents, it would be farcical because

³⁰¹ John P. Carlin, *Detect, Disrupt, Deter: A Whole-of-Government Approach to National Security Cyber Threats*, 7 HARV. NAT'L SEC. J. 391, 402-403 (2016).

³⁰² Allen, *supra* note 6, at 3.

³⁰³ Stephan De Spiegeleire, et al., *Artificial Intelligence and the Future of Defense: Strategic Implications for Small – and Medium-Sized Force Providers*, The Hague Centre for Strategic Studies 79 (2017).

³⁰⁴ Shixun You, et al., *supra* note 102, at 37447.

³⁰⁵ Michael S. Chase, PLA Rocket Force Modernization and China's Military Reforms, Testimony presented before the U.S.-China Economic and Security Review Commission on February 15, 2018, RAND Corporation (2018), <https://www.rand.org/pubs/testimonies/CT489.html>.

³⁰⁶ Allen, *supra* note 6, at 6.

³⁰⁷ *Id.*

³⁰⁸ *Id.* at 21.

³⁰⁹ *Id.* at 10.

the Chinese would just use them as a recipe book.”³¹⁰ Notably, none of the most popular machine learning software frameworks have been developed in China.³¹¹ However, China’s behavior of aggressively developing, utilizing, and exporting increasingly autonomous robotic weapons and surveillance AI technology runs counter to China’s stated goals of avoiding an AI arms race.³¹²

B. RUSSIA

Vladimir Putin announced Russia’s commitment to AI technologies stating, “Whoever becomes the leader in this field will rule the world.”³¹³ Further, Russia continues displaying a steady commitment to developing and deploying a wide range of AI military weapons.³¹⁴ In fact, Russia is significantly expanding its budget in AI cybersecurity to sway public and political opinion around the world.³¹⁵ For example, the IRA and GRU continue hacking operations relating to United States elections.³¹⁶ These efforts largely reflect effective and extensive use of AI drive cybersecurity technologies.³¹⁷

Indeed, Russia stands out as a renewed threat in cyberspace.³¹⁸ Russia has demonstrated consistent and effective capabilities in implementing AI for behavior influencing.³¹⁹ Prior to the 2016 presidential election, the IRA utilized Facebook and YouTube, targeting millions of users with advertisements aimed at influencing the election’s outcome.³²⁰ Further, in October 2017,

³¹⁰ Chris Anderson, *Elon Musk’s Mission to Mars*, Wired Magazine (October 21, 2012), <https://www.wired.com/2012/10/ff-elon-musk-qa/>.

³¹¹ Allen, *supra* note 6, at 12.

³¹² Allen, *supra* note 6, at 7.

³¹³ Congressional Research Service, *supra* note 7, at 1.

³¹⁴ Stephan De Spiegeleire, et al., *Artificial Intelligence and the Future of Defense: Strategic Implications for Small – and Medium-Sized Force Providers*, The Hague Centre for Strategic Studies 81 (2017).

³¹⁵ Kilovaty, *supra* note 224, at 158.

³¹⁶ U.S. Department of Justice, *supra* note 171, at 4.

³¹⁷ Congressional Research Service, *supra* note 7, at 24.

³¹⁸ Jon M. Garon, *Cyber-World War III: Origins*, 7 J.L. & CYBER WARFARE 1, 4 (2018).

³¹⁹ Stephan De Spiegeleire, et al., *Artificial Intelligence and the Future of Defense: Strategic Implications for Small – and Medium-Sized Force Providers*, The Hague Centre for Strategic Studies 67 (2017).

³²⁰ Jon M. Garon, *Cyber-World War III: Origins*, 7 J.L. & CYBER WARFARE 1, 8-9 (2018).

news broke of a Russian spy campaign targeted at key United States officials beginning in 2015 and lasting until the intrusion was discovered by the United States in September 2017.³²¹

In addition, Russia is establishing a number of organizations devoted to the development of military AI applications.³²² Indeed, the Russian military has been researching and developing AI robotics control systems, with an emphasis on autonomous vehicles and planes with autonomous target identification and engagement capabilities.³²³ And, in March 2018, Russia released plans for a National Center for Artificial Intelligence, among other defense related initiatives.³²⁴ Despite Russia's aspirations, some analysts argue that it may be difficult for Russia to make significant progress in AI development due to lack of funding.³²⁵ However, others argue despite trailing behind the United States and China in military funding, Russia has still managed to become a powerful force in cyberspace.³²⁶ For example, in 2013 Russia was confident enough to grant the infamous Edward Snowden political asylum against pressure from the United States.³²⁷

C. UNITED STATES

On February 11, 2019, President Trump issued an executive order aimed at establishing America's place as the global leader in artificial intelligence technology.³²⁸ The Executive Order on Maintaining American Leadership in Artificial Intelligence (Executive Order), explains the

³²¹ Jon M. Garon, *Cyber-World War III: Origins*, 7 J.L. & CYBER WARFARE 1, 6-7 (2018).

³²² Stephan De Spiegeleire, et al., *Artificial Intelligence and the Future of Defense: Strategic Implications for Small – and Medium-Sized Force Providers*, The Hague Centre for Strategic Studies 81-82 (2017).

³²³ Congressional Research Service, *supra* note 7, at 23.

³²⁴ *Id.*

³²⁵ *Id.* at 24.

³²⁶ *Id.*

³²⁷ David D. Cole, *Assessing the Leakers: Criminals or Heroes?*, 8 J. NAT'L SECURITY L. & POL'Y 107, 107 (2015). *See also* Jacob Stafford, *Gimme Shelter: International Political Asylum in The Information Age*, 47 VAND. J. TRANSNAT'L L. 1167, 1170 (2014).

³²⁸ Donald J. Trump, Executive Order on Maintaining American Leadership in Artificial Intelligence, Exec. Order No. 13,859, 84 Fed. Reg. 3967 (February 14, 2019).

United States' policy to enhance scientific, technological, and economic leadership in AI research and development guided by five principles:³²⁹

1. The United States must drive technological breakthroughs in AI across the Federal Government, industry, and academia in order to promote scientific discovery, economic competitiveness, and national security.
2. The United States must drive development of appropriate technical standards and reduce barriers to the safe testing and deployment of AI technologies in order to enable the creation of new AI-related industries and the adoption of AI by today's industries.
3. The United States must train current and future generations of American workers with the skills to develop and apply AI technologies to prepare them for today's economy and jobs of the future.
4. The United States must foster public trust and confidence in AI technologies and protect civil liberties, privacy, and American values in their application in order to fully realize the potential of AI technologies for the American people.
5. The United States must promote an international environment that supports American AI research and innovation and opens markets for American AI industries, while protecting our technological advantage in AI and protecting our critical AI technologies from acquisition by strategic competitors and adversarial nations.³³⁰

While, the Executive Order is a nice gesture in supporting development in the right direction, a clear course of action is lacking.³³¹ The United States Government has a limited role in scientific progress and development. Specifically, the only real role played in the development of technology comes from the power of the purse.³³² And, the Executive Order does not provide for new research funds.³³³

Some AI & Law scholars argue AI should be regulated by a Government agency.³³⁴ For example, Matthew Scherer argues that the starting point for regulating AI should be a statute that establishes the general principles of AI regulation.³³⁵ Scherer proposes the Artificial Intelligence

³²⁹ *Id.*

³³⁰ *Id.*

³³¹ Winston Luo, President Trump Issues Executive Order to Maintain American Leadership in Artificial Intelligence, Harv. J. J.L. & Tech. Reports, (March 06, 2019), <https://jolt.law.harvard.edu/digest/president-trump-issues-executive-order-to-maintain-american-leadership-in-artificial-intelligence>.

³³² Kate Stith, Congress' Power of The Purse, 97 YALE L.J. 1343, 1344 (1988).

³³³ Luo, *supra* note 331.

³³⁴ Scherer, *supra* note 8, at 394.

³³⁵ *Id.*

Development Act (“AIDA”), which would create an agency tasked with certifying the safety of AI systems.³³⁶ The main idea is that AIDA would delegate the substantive task of assessing the safety of AI systems to an independent agency staffed by specialists, thus insulating decisions about the safety of specific AI systems from the pressures exerted by electoral politics.³³⁷ But, it is unlikely that standard command and control models of regulation would be effective to regulate AI.³³⁸

Further, Government agencies are notorious for over-spending and political corruption, specifically in defense procurement and regulation.³³⁹ Indeed, in the words of the late John McCain, “Our broken defense acquisition system is a clear and present danger to the national security of the United States.”³⁴⁰ Despite calls for change, the military industrial complex is far too politically powerful to allow the system to improve.³⁴¹ Indeed, despite outspending Russia and China combined on defense, the United States is still falling behind.³⁴² The reason is largely attributable to billions in administrative waste and a lack of agency accountability.³⁴³ In fact, one report suggests the Chinese are confident the United States will fail to innovate, continuing to over spend maintaining and upgrading outdated systems.³⁴⁴ Others argue, no matter the potential

³³⁶ *Id.* at 393.

³³⁷ *Id.*

³³⁸ Michael Guihot, et. al., *Nudging Robots: Innovative Solutions to Regulate Artificial Intelligence*, 20 VAND. J. ENT. & TECH. L. 385, 415 (2017).

³³⁹ Whitlock, Woodward, *supra* note 212. See also *University Research Company, LLC*, 2004 WL 2496439, at 10 (Comp. Gen. Oct. 28, 2004). See also *Femme Comp Inc. v. United States*, 83 Fed. Cl. 704, 767 (2008).

³⁴⁰ Press Release, United States Committee on Armed Services, Senate Armed Services Committee Completes Markup of National Defense Authorization Act for Fiscal Year 2016, <https://www.armed-services.senate.gov/press-releases/senate-armed-services-committee-completes-markup-of-national-defense-authorization-act-for-fiscal-year-2016>.

³⁴¹ Brian S. Haney, *Automated Source Selection & FAR Compliance*, 48 PUB. CONT. L.J. (2019) (Forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3261360.

³⁴² Niall McCarthy, *The Top 15 Countries for Military Expenditure in 2016*, *Forbes* (April 21, 2017), <https://www.forbes.com/sites/niallmccarthy/2017/04/24/the-top-15-countries-for-military-expenditure-in-2016-infographic/#6ef65a0b43f3>.

³⁴³ Whitlock, Woodward, *supra* note 212.

³⁴⁴ Allen, *supra* note 6, at 8.

for AI, the Government should handle development carefully.³⁴⁵ But, the United States Government may have a more limited role in AI development than many suspect.

Private companies are driving progress in AI. For example, Google has a massive intelligence portfolio.³⁴⁶ Some argue, Google's AI technologies are scalable to an AGI model.³⁴⁷ Commercial AI products are already heavily deployed in marketing. Indeed, to take advantage of the services offered by today's major online corporation such as Google, Facebook, and Twitter, consumers are forced to give away a great deal of personal information.³⁴⁸ A person's browser history and buying habits, together with their personal information, are enough for machine learning algorithms to predict what they'll buy and how much they'll pay for it.³⁴⁹

Interestingly, Judge Baker argues, national security law serves three purposes, providing essential values, process, and the substantive authority to act, as well as the left and right boundaries of action.³⁵⁰ However, law is characterized by the relationship between a sovereign and subject acting in a habit of obedience.³⁵¹ Whether, technology companies like Facebook, Google, Amazon, Microsoft, and Apple have more sovereignty than the United States an interesting debate. Further, most AI research advances are occurring in the private sector, where talent and funding exceeds the United States Government.³⁵² As a result, militaries and

³⁴⁵ Michael Guihot, et. al., *supra* note 338, at 454.

³⁴⁶ US 2015/0100530, Methods and Apparatus for Reinforcement Learning, to Mnih, et. al., Google (2015). WO 2018/083532, World Intellectual Property Organization, Training action selection using neural networks, to Wang Ziyu, et. al., DeepMind (2016). WO 2018/083667 World Intellectual Property Organization Reinforcement Learning systems, to Silver, et.al. DeepMind (2016). WO 2018071392, World Intellectual Property Organization, Neural networks for selecting actions to be performed by a robotic agent, to Pascanu, Razvan, et.al., DeepMind (2016). WO 2018/081089, World Intellectual Property Organization, Processing sequences using neural networks, to Van Den Oord, et.al., DeepMind (2016).

³⁴⁷ Iuliia Kotseruba, John K. Tsotsos, A Review of 40 Years in Cognitive Architectures Research Core Cognitive Abilities and Practical Application, 7 (2018), <https://arxiv.org/abs/1610.08602v3>.

³⁴⁸ MURRAY SHANAHAN, THE TECHNOLOGICAL SINGULARITY 170 (2015).

³⁴⁹ *Id.*

³⁵⁰ Baker, *supra* note 12, at 6.

³⁵¹ Hart, *supra* note 293, at 50.

³⁵² Allen, Chan, *supra* note 215, at 1.

intelligence agencies depend on the private sector for essential goods and services.³⁵³ Thus, some suggest the challenges of regulating fast-moving technology are so great that industry self-regulatory approaches are often presented as the most effective mechanism to manage risk.³⁵⁴ Therefore, one argument is the United States' national security law is in the hands of private companies, rather than the Government.

Some argue the United States' technological superiority is increasingly being challenged by competitors.³⁵⁵ In truth, United States' technological superiority has already been surpassed, if not by China, certainly by the private sector.³⁵⁶ Indeed, a serious question exists as to whether the AI arms race is between governments or private firms. Matthew Scherer argues Microsoft, Google, Facebook, Amazon, and Baidu are in a private AI arms race.³⁵⁷ Of particular interest is Microsoft's recent investment OpenAI, whose stated mission is "to ensure that artificial general intelligence (AGI)—by which we mean highly autonomous systems that outperform humans at most economically valuable work— benefits all of humanity."³⁵⁸ Indeed, OpenAI, who employees some of the best AI researchers in the world recently received \$1 billion in funding from Microsoft, which may make the company's mission more difficult to differentiate from profit driven motives.³⁵⁹ Regardless, with its investment in OpenAI, Microsoft is inarguably a

³⁵³ Johnathan Wakely, Andrew Indorf, Managing National Security Risk in an Open Economy: Reforming the Committee on Foreign Investment in the United States, 9 Harv. Nat'l Sec. J. 1, 4 (2018).

³⁵⁴ Michael Guihot, et. al., *supra* note 338, at 431.

³⁵⁵ Kadtko, Wharton, *supra* note 2, at 1.

³⁵⁶ US 2015/0100530, Methods and Apparatus for Reinforcement Learning, to Mnih, et. al., Google (2015). WO 2018/083532, World Intellectual Property Organization, Training action selection using neural networks, to Wang Ziyu, et. al., DeepMind (2016). WO 2018/083667 World Intellectual Property Organization Reinforcement Learning systems, to Silver, et.al. DeepMind (2016). WO 2018071392, World Intellectual Property Organization, Neural networks for selecting actions to be performed by a robotic agent, to Pascanu, Razvan, et.al., DeepMind (2016). WO 2018/081089, World Intellectual Property Organization, Processing sequences using neural networks, to Van Den Oord, et.al., DeepMind (2016).

³⁵⁷ Scherer, *supra* note 8, at 354.

³⁵⁸ OpenAI Charter, (April 9, 2018), <https://openai.com/charter/>.

³⁵⁹ Stephen Nellis, Microsoft to Invest \$1 Billion in OpenAI, Reuters (July 22, 2019), <https://www.reuters.com/article/us-microsoft-openai/microsoft-to-invest-1-billion-in-openai-idUSKCN1UH1H9>.

major player in the future of AI development. Similarly Google's AI principles include a mission to create AI, which is socially beneficial.³⁶⁰ Yet, despite being a world leader in AI, Google's AI is used mainly for advertising, where it derives ninety-five percent of its revenue.³⁶¹

There is little social benefit or societal good to come from AI. Some promise benefits in fields like law and medicine, but AI innovation fails to solve the profit driven access problems foundational to these industries.³⁶² At a deeper level, inequality and injustice are largely supported by societal structures, blind to technological developments and AI will only exacerbate these problems. Indeed, AI is developing in corporations whose principle purpose is to maximize shareholder wealth.³⁶³ Further, the AI & Ethics school of thought is largely idealistic and academic.³⁶⁴ In a perfect world, corporate ethics would support AI development in compliance with certain principles.³⁶⁵ In reality, the United States Government has little control over profit driven big technology corporations and lacks meaningful insight into AI research.³⁶⁶ In fact, the Russian and Chinese Government also lack control over big technology corporations, relying on their research to develop their own AI systems.³⁶⁷ In sum, the dynamics of United States AI national security policy largely revolve around decisions made by corporate actors, specifically: Amazon, Google, Facebook, Microsoft, and Apple.³⁶⁸ There is an improbable, but

³⁶⁰ Sundar Pichai, AI at Google: Our Principles, (June 7, 2018), <https://ai.google/principles/>.

³⁶¹ GILDER, *supra* note 79, at 37.

³⁶² Hemant Taneja, Unscaled: How AI and New Generation of Upstarts Are Creating the Economy of the Future, 73 (2018). *See also* Kevin D. Ashley, Artificial Intelligence and Legal Analytics 3 (2018).

³⁶³ Julian Velasco, The Fundamental Rights of The Shareholder, 40 U.C. DAVIS L. REV. 407, 409 (2006).

³⁶⁴ Karl Manheim, Lyric Kaplan, *Artificial Intelligence: Risks to Privacy and Democracy*, 21 YALE J. L. & TECH. 106, 160 (2019).

³⁶⁵ Veronica Root, *Coordinating Compliance Incentives*, 102 CORNELL L. REV. 1003, 1051 (2017).

³⁶⁶ Veronica Root, *The Compliance Process*, 94 IND. L.J. 203, 231 (2019). *See also* *Elon Musk at the National Governors Association 2017 Summer Meeting*, C-SPAN (July 15, 2017), <https://www.c-span.org/video/?431119-6/elon-musk-addresses-nga>. (Musk responding to Arizona Governor Doug Ducey at 57:00-60:00).

³⁶⁷ Allen, *supra* note 6, at 12.

³⁶⁸ Alexander Tsisis, *Marketplace of Ideas, Privacy, and The Digital Audience*, 94 NOTRE DAME L. REV. 1585, 1589 (2019).

possible exception a breakthrough in AI occurs by a smaller team or single person producing AGI.³⁶⁹

CONCLUSION

Conventional wisdom teaches technological progress is driven by the Law of Accelerating Returns (LOAR).³⁷⁰ The LOAR's application to information technology, Moore's Law, projects exponential trends in technological progress toward an ultimate technological singularity.³⁷¹ This notion has developed into a school of thought called Technological Utopianism.³⁷² Technological Utopianism refers to the idea that digital life is the natural and desirable next step in the cosmic evolution of humanity, which will be good.³⁷³ As a result of Technological Utopianism, a majority of literature on the subject of technology is inherently optimistic, both in terms of outcomes and rates of progress.³⁷⁴ Yet, it is critical to resist the temptation to accept the claims of this literature.³⁷⁵ The future does not happen on its own and AI technologies could certainly have terrible outcomes.³⁷⁶

One argument for the future of the United States Government in AI development is to pursue an open government model. Open government is a concept referring to the free flow of information between the Government and the public.³⁷⁷ The goal of such a model would be to improve transparency, education and access to critical AI information.³⁷⁸ As a result, AI issues could be discussed, debated, and decided democratically. However, in practice there is little hope

³⁶⁹ BOSTROM, *supra* note 11, at 101.

³⁷⁰ Haney, *supra* note 22 at 155.

³⁷¹ KURZWEIL, *supra* note 9, at 250.

³⁷² TEGMARK, *supra* note 18, at 32.

³⁷³ MARTINE ROTHBLATT, VIRTUALLY HUMAN 283 (2104).

³⁷⁴ BOSTROM, *supra* note 11, at 34.

³⁷⁵ Peter Thiel, The Education of a Libertarian. Cato Unbound (May 1, 2009).

³⁷⁶ PETER THIEL, ZERO TO ONE 195 (2014).

³⁷⁷ Mark Fenster, *The Opacity of Transparency*, 91 IOWA L. REV. 885, 895 (2006).

³⁷⁸ Joshua Apfelroth, *The Open Government Act: A Proposed Bill to Ensure the Efficient Implementation of The Freedom of Information Act*, 58 ADMIN. L. REV. 219, 220 (2006).

such a model would be put into practice. This is particularly true in the United States where agencies fight tooth and nail to hide information to which the public has a right via FOIA litigation.³⁷⁹

A second argument is that AI technology's likely dissemination into the wrong hands resolves the Fermi Paradox. A paradox is a set of arguments with apparently true propositions, leading to a false conclusion.³⁸⁰ Consider, the Milky Way is one of hundreds of billions of galaxies in the Universe, each containing hundreds of billions of stars.³⁸¹ Commonly, these stars contain Earth-like planets.³⁸² As a result, statistically it is almost certain life would have developed somewhere else in the Universe before life on Earth.³⁸³ And yet, mankind finds itself bound to a pale blue dot on the outskirts of the Milky Way, apparently alone in the Universe. Fermi's Paradox asks the question, "Where are they?"³⁸⁴

The great British Mathematician Irving J. Good argued AGI would be the "last invention that man need ever make."³⁸⁵ And the late Stephen Hawking observed, "The development of artificial intelligence could spell the end of the human race."³⁸⁶ Further, both Nick Bostrom and Max Tegmark have argued persuasively, humans may not be able to control AGI.³⁸⁷ These observations provide support that AI may lead to a catastrophic event resolving the Fermi Paradox.

³⁷⁹ John C. Brinkerhoff Jr., FOIA's Common Law, 36 Yale J. on Reg. 575, 576 (2019). (FOIA is an acronym for Freedom of Information Act.)

³⁸⁰ MARGARET CUONZO, PARADOX 2 (2014).

³⁸¹ CARL SAGAN, PALE BLUE DOT A VISION OF THE HUMAN FUTURE IN SPACE 21 (1994).

³⁸² Nick Bostrom, *In the Great Silence there is Great Hope 2* (2007) <https://nickbostrom.com/papers/fermi.pdf>.

³⁸³ *Id.*

³⁸⁴ *Id.* at 5.

³⁸⁵ Irvin J. Good, *Speculations Concerning the First Ultraintelligent Machine*, 6 ADVANCES IN COMPUTERS 31 (1965).

³⁸⁶ Rory Cellan-Jones, *Stephen Hawking Warns Artificial Intelligence Could End Mankind*, BBC News (Dec. 2, 2014), <http://christusliberat.org/wp-content/uploads/2017/10/Stephen-Hawking-warns-artificial-intelligence-could-end-mankind-BBC-News.pdf>.

³⁸⁷ BOSTROM, *supra* note 11, at 155. See also TEGMARK, *supra* note 18, at 176.

In sum, the United States' national security is now dependent, not on its Military or Defense Agencies, but on big technology companies. In part because big technology companies have powerful influence over political decision makers.³⁸⁸ Further, big technology companies have the most talented people and own the rights to the most powerful weapons. Yet, the answer is not to break up big technology companies, which disadvantages the United States compared to our adversaries.³⁸⁹ Instead the answer is to accept the changing power dynamics and do the best we can with a broken political system. The only alternative would be revolution.³⁹⁰

³⁸⁸ Megan Henney, Big tech has spent \$582M lobbying Congress. Here's where that money went, FOXBusiness (July 26, 2019), <https://www.foxbusiness.com/technology/amazon-apple-facebook-google-microsoft-lobbying-congress>.

³⁸⁹ Sheelah Kolhatkar, How Elizabeth Warren Came Up with a Plan to Break Up Big Tech the New Yorker (August 20, 2019), <https://www.newyorker.com/business/currency/how-elizabeth-warren-came-up-with-a-plan-to-break-up-big-tech>.

³⁹⁰ National Archives, The Declaration of Independence: A Transcription (July 4, 1776), <https://www.archives.gov/founding-docs/declaration-transcript>.

APPENDIX A. SUMMARY OF NOTATION

<i>Summary of Notation</i>	
Notation	Meaning
π^*	Optimal policy.
$Q(s, a)$	Q-function.
(s, a)	State-action pair.
ϕ	Q-function parameters.
γ	Discount factor.
$\mathbb{E}[x]$	Expectation of random variable.
$\arg \max_a f(a)$	A value of a , at which $f(a)$ takes its maximal value.
r	Reward.
θ_k	Policy parameters for k experiment.
$L(s, a, \theta_k, \theta)$	Objective function.
$A^{\pi_{\theta_k}}$	Advantage estimate for policy given parameters.
$\pi_{\theta}(a s)$	The policy given parameters.
ϵ	Hyperparameter defining how far away the new policy is allowed to go from the old.
$a^*(s)$	Optimal action-value function.
\mathcal{D}	Replay Buffer.

APPENDIX B. SIGNIFICANT CYBER INCIDENTS DATA

Year	China	Russia
2009 (September 2008 - August 2009)	3	2
2010 (September 2009 - August 2010)	9	3
2011 (September 2010 - August 2011)	4	0
2012 (September 2011 - August 2012)	4	0
2013 (September 2012 - August 2013)	8	4
2014 (September 2013 - August 2014)	3	0
2015 (September 2014 - August 2015)	7	5
2016 (September 2015 - August 2016)	3	5
2017 (September 2016 - August 2017)	4	11
2018 (September 2017 - August 2018)	13	27
2019 (September 2018 - August 2019)	24	30