Dylan Levsey

ENG

Spring 2017

The role of Government Regulation involving Artificial Intelligence

Despite what the movies portray, robots and other artificial intelligent technologies could ultimately advance the human race. Einstein once said, “[w]e must not condemn man because his inventiveness and patient conquest of the forces of nature are being exploited for false and destructive purposes. Rather, we should remember that the fate of mankind hinges entirely upon man’s moral development…” If we want to avoid the pit falls of artificial intelligence, we should not focus on restricting development—we should be encouraging the goodwill of developers.

Artist and scientist both romanticize and fear the idea of artificial intelligence. Some believe that artificial intelligence will be the end of civilization. Some even go as far as depicting these fears as indestructible time traveling robots (*Terminator*). In 2014, Theoretical Physicist Stephen Hawking told the BBC that “[t]he development of full artificial intelligence could spell the end of the human race.” However, evidence suggest otherwise.

AI could transform the way we complete tasks and how fast we are able to move on to another one. With AI, going to other planets would be a simple task; Cancer treatments could become as simple as getting a vaccination. The conservation of resources would become easier. Surgeries that are invasive and dangerous could

become mundane and simple; the concept of life expectancy, or death, could become outdated.

If the government wants to avoid potential risk of AI, they should stray from stringent regulations, or rules, that can stifle innovation; innovations that can be used for the greater good. Governmental regulations should not be imposed on AI development based upon fear, as there is no reason to fear AI. Stringent regulations would also prevent companies, small or big, that are developing technologies from expanding or even existing, as well as inventing technologies that would benefit the world. The tight regulations would also be counterintuitive as the potential growth for AI is highly disputed and unknown. Slowing the research in development would not only affect the country at which the government resides, it would be in the countries best interest not to; because while AI is not inherently evil, even with global monitoring, there is no way to predict another countries intentions or pace of development.

To understand the significance of an artificially intelligent world, we should first understand some relevant terminology. The first term we should understand is artificial intelligence. Does it mean that we could make technology smarter than humans or maybe it means that technology could make humans smarter? To be artificially intelligent, does a machine have to be self-aware or can it simulate self-awareness? The website for *Oxford Dictionary* describes it this way: “The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision making, and translation between languages” (*Oxford). Funk & Wagnall’s New World Encyclopedia*’s version describes

artificial intelligence with a more humanistic approach AI:

> Artificial intelligence (AI), a term that in its broadest sense would indicate the ability of an artifact to perform the same kinds of functions that characterize human thought processes... With the growth of modern science, the search for AI has taken two major directions: psychological and physiological research into the

> nature of human thought and cognitive processes, and the technological development of increasingly sophisticated computing systems. [These] simulate various aspects of such human activities as reasoning, learning, and self
>
> correction and that can do tasks traditionally considered to require intelligence, such as speaking and speech recognition, vision, chess playing, and disease diagnosis. (“Artificial Intelligence”)

The definition of AI is slightly more flexible than some might believe. The *Oxford Dictionary* seems to describe it as something that may have human-type intellect; while *Funk & Wagnall’s New World Encyclopedia*’s version describes it as something that merely simulates human-type intellect. The two definitions may be slightly different. Nevertheless, they display the slight variations in the understanding of what AI can mean. However slight they may be.

Artificial intelligence is all around us: Amazons shopping recommendations, Facebook Image Recognition, Netflix streaming video, Siri, Cortana, Alexa, and other voice assistant and photo augmentation apps. Although these examples are all far

away from being confused for human intelligence, they are very sophisticated. For instance, Facebooks Image Recognition can identify a person 98 percent of the time. To put that into perspective, the FBI's facial recognition system "identifies the correct person in the list of the top 50 people only 85 percent of the time." The most interesting aspect of Facebooks Image Recognition is that "[e]very time one of its 1.65 billion users uploads a photo to Facebook and tags someone, that person is helping the facial recognition algorithm." Facebook's image recognition is reasoning out who is in a photo, learning how to get better at it, and correcting itself to better identify a person: Facebook is using artificial intelligence (Intel; LaChance).

Facial recognition tools like the one Facebook uses are constantly under scrutiny for privacy concerns. However, recently there has been even more attention to the safety of the self driving car. The self-driving car has long been a fantasy of many people. Now this seemingly unrealistic idea has become seemingly realistic. Self-driving cars serve several good purposes: they can transport the visually impaired, the disabled, and make it safer for anyone who can drive.

The self-driving car operates on a very sophisticated system. While every vehicle is designed differently, how the self-driving car operates, what they are equipped with, and how functional they are seem to remain constant. For example, the majority of self-driving cars include the following: multiple cameras, Radar sensors, a LIDAR unit (or laser unit), and a computer. The cameras have the ability to capture multiple images making it the main component for navigation; they can detect stationary objects like traffic signals and stop signs, or moving objects like pedestrians and bicyclists. The

LIDAR unit uses laser beams to further the cars vision accuracy; the unit is able to capture a “360-degree image of the car’s surroundings.” The radar sensors will pick up the rest of the slack by measuring the distance between the vehicle and the obstacles surrounding it. The onboard computer collects this information, analyzes it, and compares it to the maps already stored in its system. Not all autonomous, or self-driving, vehicles are the same though. For example, Tesla’s version of the autonomous vehicle uses a vision system in lieu of the LIDAR technology seen on most self-driving cars. Meaning Tesla’s version is only self-driving on the highway. The benefits of both types of systems are vast (Gates, et al.).

A key benefit of the autonomous car is that it is sophisticated enough to understand if a different car is in the wrong lane. The sophistication continues to surmount when the driver has

the ability to set the cruise control to a threshold based on distance rather than a speed. For example, the cruise control can be set to a minimum of the three car lengths away from surrounding vehicles rather than a set speed limit of 60-mph. The cameras and radar allow the car to detect blind spots, and self-park (with minor assistance). The radars, lasers, and cameras coordinate well enough to have the ability to predict impending collisions. This level of sophistication means that the car itself can decide to slow down before anyone even knows the conditions are dangerous Some of these capabilities are captured in a video that surfaced late in 2016. In the video, the Tesla-car notices that there is a high-possibility of getting into a wreck. The car begins braking as it warns the driver (the warning sound is a loud

beep) and within seconds, there is a collision up ahead. The Tesla vehicle remains unscathed (Gates, et al.; McCausland & Walters).

The scrutiny comes into play when the self-driving car no longer has to just slow down, but make a decision that would have dire consequences. Computers govern these vehicles. Algorithms govern these computers, and people govern these algorithms. In the event of an accident, the algorithm tells the computer exactly what to do. This makes the car insanely quick and equally as precise. In addition, the computers avoid freaking out and making deadly mistakes because of their lack of emotion. The sophistication of the computer means that in the event of a car wreck the autonomous vehicle will know exactly what to do, to get the result the algorithm tells it to have. The most common result being the passenger of the self-driving car to remain alive and well. However, if the impending wreck involves an autonomous vehicle with one passenger, and another vehicle with five passengers, what should the autonomous car do? Does the car save the family and sacrifice the driver; or does the car save the driver and sacrifice the family? While certain aspects of self-driving cars seem advantageous, it brings up an ethical

dilemma. If programming the car creates a situation in which the value of one's life is measured, what decides that value? Also, who gets to decide? (Nyholm & Smids).

Ethics described by *Webster's Online Dictionary is* "the discipline dealing with what is good and bad and with moral duty and obligation." We deal with ethical decisions every day; ethical decisions based on the ethical values of the culture we grow up in. How would ethical decisions work with AI? Consider this: the billionaire

entrepreneur Elon Musk just started investing in a company “that wants to develop ways for human brains to interact with software” or develop “’neural lace’ technology that would allow human brains to interface with machines without a physical connection” The company’s name is Neuralink, and it is on a mission to change the way we learn, compete, and remember information. While the potential for the company Neuralink is looking good, the way they could reach that potential is unclear. (Gates, et al.; Edelstein).

Neural lace technology is classified as a brain-computer interface (BCI); these medical technologies aim to “improve the quality of life of people who are paralyzed or have severe motor disabilities.” BCI experimentation would not be begin until 1969; the field of study regarding BCI’s would not emerge until the 1970’s; and the invention of neuro-lace technology would not happen for another 45 years. In 2015, the experiment done on neuro-lace technology included inserting some very thin electronic-mesh into the brain tissue of mice. Connecting a wire from the mesh to a computer, they were able to tell if the brain cells were still able to grow. The results from the experiment found that the mice were not only growing brain cells around the mesh; they were growing brain cells because of the mesh (Wu et al.; Valeriani; Newitz).

The progress of the brain-computer interface technologies has proven to be slow, but substantial. In areas such as the spine and the brain, they were able to make leaps such as,

“strengthen[ing] connections between two brain regions or between the brain and the spinal cord.” This allowed them to “reroute information to a paralyzed limb.” They have

recently had success using similar types of treatment for diabetes; however, there are negative aspects surrounding the BCI procedures. The BCI's procedures are currently quite invasive and require a surgical procedure of inserting electrodes into the brain, raising the risk of brain damage. There are less risky procedures involving BCIs that do not require invasive procedures e.g., magnetic pulses; however, they are not as prominent or proven (Wu et al.).

The ultimate goal of the BCI future would be to acquire the ability to think quicker, know more with less time, and to be able to solve problems that are more abstract. To do this there would need to be a wireless connection of the brain to the cloud system. Although the idea of having the brain connected to a cloud system sounds unsafe, some say it is the next natural step:

> Connecting our brains directly to technology may ultimately be a natural progression of how humans have augmented themselves with technology over the ages, from using wheels to overcome our bipedal limitations to making notations on clay tablets and paper to augment our memories. (Wu et al.)

The drive for BCIs stems from the belief that we will need new ways to compete with the AI of tomorrow. Elon Musk has gone as far as saying that we need advancement in BCIs so that we do not become housecats to the Artificially Intelligent (D'Onfro).

With the brain connected to a wireless network, would there be some privacy concerns? Would the neuro-lace implants make us smart enough to be able to prevent malicious attacks? Perhaps identity theft would have a whole meaning. For instance,

identity theft could be the switching out of information between peoples brains. However, given the amount of benefits that would stem from advancing BCI's, this scenario is most likely not going to happen.

With everyone connected to a brain enhancing device and having other AI technologies alongside it, people may not feel the need to be malicious. The higher level of thinking would help us eradicate diseases, hunger, and inequality. With the promising features of BCIs and their ability to improve people's lives from various diseases and disabilities: being paralyzed, blind, strokes, spinal injuries, brain injuries, or those who have had severe motor disabilities. The reality of a neuro-laced, or neuro-technological, society may not be too far away. It all depends on the growth (Wu et al.).

Seen as the scariest part about artificial intelligence; or seen as the most promising aspect. The technological growth for AI is very important to understand. AI has the potential to grow exponentially, or as one author put it:

> AI could duplicate itself very quickly, in no more time than it takes to make a copy of the AI's software. Moreover, since the AI copy would be identical to the original, it would be born completely mature, and the copy could begin making its own copies immediately. Absent hardware limitations, a population of AIs could therefore grow exponentially at an extremely rapid rate, with a doubling time on the order of minutes or hours rather than decades or centuries. (Bostrom & Yudkowsky)

It is important to understand just how fast AI grows. The possibility of

exponential growth is the main reason some people may fear the technology. While maliciously used technology could grow exponentially, the beneficial aspects of AI could grow exponentially as well. The rapid reproduction of artificial intelligence began in 1936 and not with a robot, but with a human (“Alan Turing").

Although the term “artificial intelligence” had not been coined until 1955 by computer Scientist John McCarthy, the history of AI really begins with mathematician Alan Turing. Turing was born in London, on June 23,1912, and in 1936 & 1950 he would publish papers that would build the foundation for artificial intelligence. Alan Turing's paper in 1936, “On Computable Numbers, with an Application to the Entschediungsproblem [A decision problem],” was the main concept the modern computer is based on and was the first paper to present the idea of a universal machine. The paper in 1950 posed the question, “Can machines think?” and he started out this paper with six philosophical driven answers to this question. The first section of his paper is where we get the term *Turing test*. Author Stephan Muggleton, a member of the Department of Computing at the Imperial College in London, UK summarizes this section (Russell;“Alan Turing"; Muggleton):

> The imitation game. Often referred to as the “Turing test”, this is a form of parlour [indoor] game involving a human interrogator who alternately questions a hidden computer and a hidden person in an attempt to distinguish the identity of the respondents. The Imitation Game is aimed at providing an objective test for deciding whether machines can think. (Muggleton)

Even today, the Turing test is brought up in various aspects of artificial intelligence: peer reviewed journals, experiments, gameshows, and especially literature. The Turing test has been tested with several different computers. These computers have played various games that require the player to think: IBM's Watson computer won against humans on the gameshow "Jeopardy!" (Markoff); in 1997 IBM computer Deep Blue beat world chess champion in 6 games ("Deep Blue"); and most recently Google's AlphaGo beat master Go player 3-0 in a best-of-five competition: "Go" is a Chinese board game that "is [considered] a much more complex

challenge for a computer than chess" (Markoff; "Deep Blue"; "Artificial intelligence: Google's AlphaGo").

The Turing test was most notably discussed in 2014 when, "[a] computer program called Eugene Goostman, which simulates a 13-year-old Ukrainian boy, is said to have passed the Turing test at an event organised by the University of Reading." The Turing test is passed if the judges think the computer is human more than 30% of the time in "a series of five-minute keyboard conversations." The experiment was put together by gathering a panel of 30 judges, five computers, and some humans. The judges would communicate with the test subjects by instant messenger in 5-minute intervals; each judge would go through 10 test. Based on the conversations, the judges would then say if they believed they were messaging a human or a computer. Eugene was able to convince the judges 33% of the time that it was human, and it only had to convince the judges 30 % of the time it was human to pass the Turing test ("Computer AI Passes Turing"; Reading.ac.uk.).

The organizers behind the test hailed it as "historic." However, some professors like Prof Noel Sharkey say it was rigged: “It was very clever ruse to pretend to be a 13-year-old Ukranian [sic.] boy, which would constrain the conversation. But these competitions are really great to push developments.” Prof Stevan Harnad stated, "It's nonsense…We have not passed the Turing test. We are not even close." Experiments designed the way Eugene’s was usually require a level of imagination reserved for writers, directors, or idealist ("Computer AI Passes Turing”).

The idea of a robot can be found in literature work as early as the 1950’s. One of those works being the book *I’Robot* by Isaac Asimov. *I’Robot* was released in the year of 1950, making Asimov truly ahead of his time—it would be almost 30 years until the moon landing— Born January 2, 1920, he would go on to write 500 books, most of which were based around

science fiction. The book *I’Robot* is credited as being the origin of the word *robotics* ("Isaac Asimov"). The book is also notorious for creating the three laws of Robotics, which are the following:

1) A robot may not injure a human being or, through inaction, allow a human being to come to harm

2) A robot must obey orders givein to it by human beings except where such orders would conflict with the First Law.

3) A robot must protect its own existence as long as such protection does not conflict with the First or Second… (Asimov)

There have been many science fiction pieces that have been influenced by Asimov's works. One of the most recent examples being a film that was released in 2015, titled *Ex Machina,* in which a human subject is selected to participate in a Turing test. The main theme of this movie demonstrates a Turing test where the robot is not disguised as a human, but is shown as a robot. The purpose of this is that if the test subject can see the robot and still end up convinced it is not a robot, then it truly passes the Turing test (*Ex Machina*).

The impact sci-fi can have on the science community can be major. Elon musk for instance is widely known for the influence sci-fi has had on him. According to CNBC Elon Musk once told the Guardian that Asimov's books helped him learn "civilizations move in cycles," and that Asimov's books helped drive his ambition. He has most recently been quoted as saying, "Reread Asimov's Foundation series…Brilliant" (Wood; Baer).

Sci-fi ends where reality sets in, reality sets in between the limitations of the technology available and the hurdles that should be made to advance that technology. The hurdle could be scientific and methodical, like an equation, or it could be subjective and political, like public

perception. Both of which are currently present in today's society. While it is true that the technology has reached a superior level, it is also true that it has not reached the all-knowing, all curing, and all wealthy characteristics commonly associated with an advanced society with fully functional AI. In the situation of developing AI technology, the more pertinent hurdle is the political aspect, as public perception of AI is

undoubtedly grim. Currently 72% of a 6,000- person survey believe they understand what AI is, 34 % believe they have interacted with AI, and they found that 84 percent of them actually interacted with AI on a day-to-day basis. Furthermore, 10% of them described AI as having the capability to take over the world. And when it comes to the public's perception surrounding AI regulation, a clear majority, 71%, believe that there should be regulations enacted. In fact, it has become a non-partisan issue, 74% of those who believed there should be regulation were republicans, while 73% of those who believed there should be regulation were democrats. The general population is paranoid and fearful over a technology that should be embraced.

The idea of a hostile AI system stems from the belief that if AI gained superintelligence without gaining empathy it would be rendered sociopathic; resulting in a super advance computing power that would have a strong desire to end all human life. However, if AI eventually becomes sentient or "Intelligent," then it must be a separate species made of completely different parts, or even require different amounts of energy. If AI is a completely different species made up of different parts, requiring different amounts/types of energy, what more we would know about how the mind of AI works compared to how the mind of a chimpanzee, who we have a 99% commonality with, or even something our DNA has a 50% commonality with such as a banana. For that matter, what incentive would a super-intelligent robot have to be evil, and how would that differ from that of the incentive of a bananas? If the AI

did not like us, why would they not use their super-intelligence to just go to a different universe? If they were super-intelligent, would they even want go to a different

universe; would they not be able to innovate enough resources for everyone, would the lack of innovation render them “not super-intelligent” but just tin versions of homeo-sapiens? Some experts suggest that it would be a matter of resource management, that in order for the Super-Intelligent beings to secure their survival, they would have to secure human extinction. However, the hostility and potential world ending scenarios of AI has been disputed. For example, Stanford University’s study on AI, dubbed the “One Hundred Year Study on Artificial Intelligence,” found that there was “no cause for concern that AI is an imminent threat to humankind.” With that same study protesting for a more open-minded society towards AI: “if society approaches AI with a more open mind, the technologies emerging from the field could profoundly transform society for the better in the coming decades.” AI being a generally a safe technology, and that while some programs could possibly run amuck, it would be more likely that the AI technology would be controlled by those with ill intent rather than the AI having ill intentions. Furthermore, the fact that AI can do so much good in the world is even more reason as to why it is not evil.

AI has the capacity to possibly save numerous lives, transform the roads, as well as save the environment, which would end up benefiting the globe. For example, the Obama Administrations Archives includes a report on potential AI systems. This report was published in October of 2016, and it is an extensively detailed description of the information the government has gathered and analyzed on technology involving AI .i.e., Information gathered from numerous studies and testimonies from researchers that are experts in field of AI. A study quoted in the report surrounding healthcare indicated that the error rates for lymph node cancer diagnoses among healthcare professionals

actually goes down when paired with an AI system. Both of

these test subjects actually do worse apart, indicating hope for a symbiotic relationship between AI and humans. Researchers also determined that “AI-based applications could improve health outcomes and the quality of life for millions of people in the coming years.” The health benefits

of AI can extend to include autonomous vehicles because they are safe, and they are highly efficient when it comes to environmental causes. For instance, researchers from both studies found that “AI-enabled smarter traffic management applications are reducing wait times, energy use, and emissions by as much as 25 percent in some places,” as well as “[t]he reduced need for parking [which] may affect the way cities and public spaces are designed.” Not only would commute and depollution be positively affected by AI enabled autonomous vehicles, but AI could also help advance preservation techniques:

> Autonomous sailboats and watercraft are already patrolling the oceans carrying sophisticated sensor instruments, collecting data on changes in Arctic ice and sensitive ocean ecosystems in operations that would be too expensive or dangerous for crewed vessels. Autonomous watercraft may
>
> be much cheaper to operate than manned ships, and may someday be used for enhanced weather prediction, climate monitoring, or policing illegal
>
> fishing. (Executive Office of the Obama Administration)

The beneficial aspects of AI are vast and can hardly go overlooked: the healthcare system would be much more accurate, roads would be more efficient, and environmental protection efforts would be made easier. The caveat being that this is all depending on how reactive the government is.

Government regulations, or government rules, can be seen as "not doing enough," or "doing too much." However, sometimes they are necessary as they can protect American citizens

and sometimes citizens of the world, and no matter how well intended regulations may be, no one ever agrees on whether they are good or not. The poll aggregating website, P*ew Research Center,* showed that 49% of Americans believed the government wasn't doing enough, and it showed that 43 percent believed that the government was doing too much. As polarizing as some

of these regulations may be, people are in near agreeance of imposing regulations on the development of AI technologies (71% believing in imposing regulations). Certain subfields of the AI community receive more forms of regulations than others. Autonomous vehicles are mostly forgotten in the area of regulation: there are only four states that have any laws regarding autonomous vehicles (Nevada, Florida, California, and Michigan). Most regulations as of right now are on the healthcare system and these are outdated. For example, there was a huge leap "in collecting useful data from personal monitoring devices and mobile apps, from electronic health medical procedures and hospital operations." This leap was then pushed backwards by the

momentum of regulation, hindering the possibilities of a better healthcare system based on data collection. There have been “incentives provided by the Affordable Care Act [that] have accelerated the penetration of electronic health records (EHRs).” The only issue here being that the execution was poor, making clinicians doubt the incentives. While the government has addressed these issues, sometimes it seems it is easier said than done. The department in charge of preparing America for jobs of the future, The Department of Labor (R&D), for example is currently under a budget of only $14 million. With departments like the R&D only able to pull in $14 million, it will be hard for certain companies to invest in AI technology that could transform America’s (and perhaps the globes), economy. The conventional wisdom when it comes to AI is that it is the ultimate job killer, and regulation should be enacted to save everyone from the destruction of our economy; this is

however the furthest thing from the truth. While it is partly true that AI will kill some jobs, Oxford for instance, did a study that found there would be as much as 47 percent of American jobs that would be automated by 2033. However, there is evidence that suggest that job loss would only be temporary. With some experts saying that “AI wont so much obliterate the job market as disrupt it.” Furthermore, a study conducted in 2013 found that the commercial drone industry alone “could generate more than $82 billion for the U.S. economy and create more than 100,000 new jobs over the next 10 years…Tax revenue in the states was predicted to increase more than $482 million in the first decade after integration.” That sector is not alone though as Stanford’s 100 year study the expansion and growth of AI will not just create new jobs, it will also make

the cost of living a lot cheaper, and that it is such a radical way for wealth creation that "everyone should be entitled to a portion of the world's AI-produced treasures." It could also breed innovation and spur a whole new way of looking at education. Newer jobs are going to require more data analysis, and less manual labor, and hopefully future generations be able focus on what they love to learn or do.

Works Cited

"Alan Turing." Biography.com. A&E Networks Television, 21 Oct. 2016. Web. 04 Apr. 2017. ARTIFICIAL INTELLIGENCE AND LIFE IN 2030 (2016): 1-52. Ai100.stanford.edu/. Stanford, Sept. 2016. Web. 28 Apr. 2017.

"Artificial Intelligence." Funk & Wagnall's New World Encyclopedia, 2016, p. 1p. 1. EBSCOhost, search.ebscohost.com/login.aspx?direct=true&db=funk&AN=AR151400&site=ehost-live.

"Artificial Intelligence." Oxford Dictionary. Oxford, Web. 2 Apr. 2017.

Artificial Intelligence: Google's AlphaGo Beats Go Master Lee Se-dol." BBC News. BBC, 12 Mar. 2016. Web. 06 Apr. 2017

Asimov, Isaac. "I, Robot." Barnes & Noble. Web. 06 Apr. 2017.

Baer, Drake. "This Apocalyptic Sci-Fi Novel Convinced Elon Musk to Make Humans An Interplanetary Species." Business Insider. Business Insider, 09 Dec. 2014. Web. 06 Apr. 2017.

Bostrom, Nick, and Eliezer Yudkowsky. "The Ethics of Artificial intelligence." The Cambridge Handbook of Artificial intelligence: 316-34.

Http://www.nickbostrom.com. Cambridge University Press, 2011. Web. 3 Apr. 2017.

Bostrom, Nick. Superintelligence: Paths, Dangers, Strategies. Oxford: Oxford UP, 2016. Print. Cellan-Jones, Rory. "Stephen Hawking Warns Artificial Intelligence Could End Mankind." BBCNews. BBC, 02 Dec. 2014. Web. 18 Apr. 2017.

“Computer AI Passes Turing Test in 'world First'." BBC News. BBC, 09 June 2014. Web. 06 Apr. 17

Couch, Christina. "All in for AI." PM Network, vol. 29, no. 10, Oct. 2015, pp. 12-13. EBSCOhost, search.ebscohost.com/login.aspx?direct=true&db=buh&AN=110038511&site=ehost-live.

“Deep Blue." Ibm.com. IBM, Web. 06 Apr. 2017.

D'Onfro, Jillian. "Elon Musk Thinks We Need Brain-computers to Avoid Becoming 'house Cats'to Artificial Intelligence." Business Insider. Business Insider, 02 June 2016. Web. 21 Apr. 2017.

Edelstein, Stephen. "Elon Musk's New Company, Neuralink, Looks to Upgrade Your Brain With AI Tech." The Drive. 29 Mar. 2017. Web. 03 Apr. 2017.

Einstein, Albert. Einstein on Politics: His Private Thoughts and Public Stands on Nationalism. Place of Publication Not Identified: Princeton U Pres, 2013. Print

“Ethics.” Webster Dictionary. Webster, Web. 11 Apr. 2017.

Ex Machina. Dir. Alex Garland. Perf. Alicia Vikander, Domhnall Gleeson, Oscar Isaac |. Universal Studios, 2015. DVD.

Examples of Artificial intelligence in Everyday Life to illustrate why some of our moral norms need to be rethought in the context of AI reproduction," Intel. Intel, Web. 02 Apr. 2017. Gates, Guilbert et, al. "When Cars Drive Themselves." The New York Times. The New York Times, 14 Dec. 2016. Web. 21 Apr. 2017.

Isaac Asimov." Biography.com. A&E Networks Television, 02 Apr. 2014. Web. 06 Apr. 2017. LaChance, Naomi. "Facebook's Facial Recognition Software Is Different From The FBI's. Here's Why." NPR. NPR, 18 May 2016. Web. 02 Apr. 2017.

Lohr, Steve. "Robots Will Take Jobs, but Not as Fast as Some Fear, New Report Says." The New York Times. The New York Times, 12 Jan. 2017. Web. 29 Apr. 2017.

Markoff, John. "Computer Wins on ‘Jeopardy!’: Trivial, It’s Not." The New York Times. The New York Times, 16 Feb. 2011. Web. 06 Apr. 2017.

McCausland, Phil and Walters, Shamar "Tesla's 'Autopilot' Begins Braking for Wreck Before Driver." NBCNews.com. NBC Universal News Group, 28 Dec. 2016. Web. 21 Apr. 2017.

Muggleton, Stephen. "Alan Turing and the Development of Artificial intelligence." AI Communications, vol. 27, no. 1, 2014, pp. 3-10. EBSCOhost, search.ebscohost.com/login.aspx?direct=true&db=psyh&AN=2013-41692- 001&site=ehost-live

Newitz, Annalee. "Scientists Just Invented the Neural Lace." Gizmodo. Gizmodo.com, 15 June 2015. Web. 21 Apr. 2017.

Nyholm, Sven, and Smids, Jilles "The Ethics of Accident-Algorithms for Self-Driving

Cars: An Applied Trolley Problem?" SpringerLink. Springer Netherlands, 28 July 2016. Web. 03 Apr. 2017.

President, Executive Office of the Obama Administration. The. PREPARING FOR THE FUTURE OF ARTIFICIAL INTELLIGENCE (2016): Obamawhitehouse.archives.gov. The US Government, Oct. 2016. Web. 24 Apr. 2017.

Russell, Mark. "Man Who Coined 'Artificial intelligence' Dead at 84." Newser. 26 Oct. 2011. Web. 04 Apr. 2017.

Scherer, Matt. "Op-ed: Poll Shows Strong Support for AI Regulation Though Respondents." Future of Life Institute. Matt Scherer Https://futureoflife.org/wp-content/uploads/2015/10/FLI_logo-1.png, 14 Apr. 2017. Web. 29 Apr. 2017.

Terminator. By James Cameron and Brad Fiedel. Perf. Arnold Schwarzenegger, Michael Biehn, and Linda Hamilton. Metro Goldwyn Mayer, 1984.

“United States DEPARTMENT OF LABOR." United States Department of Labor. The United States, Web. 29 Apr. 2017.

“University of Reading." Reading.ac.uk. University of Reading, 08 June 2014. Web. 06 Apr.2017.

Valeriani, David. Post-doctoral Researcher in Brain-Computer Interfaces, University of Essex. "Elon Musk Wants to Merge Man and Machine – Here's What He'll Need to Work out." The Conversation. 14 Apr. 2017. Web. 21 Apr. 2017

Ward, Marguerite. "Billionaire Elon Musk Credits His Success to These 8 Books." CNBC. CNBC, 06 Mar. 2017. Web. 21 Apr. 2017.

Wood, Nathaniel. "Sci-Fi Helped Inspire Elon Musk to Save the World." Wired. Conde Nast, 13 June 2015. Web. 06 Apr. 2017.

Wu, James, and Rajesh P.N. Rao. "How Close Are We to Elon Musk's Brain-computer Interface?" CNN. Cable News Network, 12 Apr. 2017. Web. 21 Apr. 2017. www.pega.com. Issue brief. PegaSystems, Web. 29 Apr. 2017. <https://www.pega.com/sites/pega.com/files/docs/2017/Apr/what-consumers-really think-about-ai.pdf>.