Bushra Ibrahim

Dr. O'Conner

HN151

31 March 2020

The Threat of Artificial Intelligence

Humans tread a fine line between saving the world and ending it, but we lack the caution that we ought to have in doing so. Artificial Intelligence (AI) teeters on this theoretical tightrope as mankind enters a new era — an era of heightened intellect, lightning-fast speeds, and unfathomable amounts of data. Stepping into this new age, we fear that the most intelligent being on Earth may be replaced, and that our own creations could leave us in the dust. Technological advancement has surged more in the past decade than humanity has in the last few centuries, and with such a massive jump comes massive consequences, and that is the case with AI.

Although contained within human hands for now, AI is accelerating at such an unprecedented rate that our so-called containment is becoming inadequate. Unchecked, the intelligence we create could turn on us, leading the way into a machinated dystopia where we have lost control. Considering most of us do not look forward to a future in which we are no longer on top, the search for a solution has begun. From this search emerges two subsequent courses of action. The first is to terminate all technological advancements entirely, an impossible feat, which leaves us only with option two. We must embed human values and modes of control into AI from the beginning of the development process, so that they will take precedence when the time comes.

The Singularity

Why is it important to hasten our security measures right now? It is because with each passing day, we inch closer to a hypothetical point in time known as the "Singularity," the last stand that all AI paths seem to lead to. At that point, a superintelligence would be born with equal, if not superior, intelligence to humans. An article by Cambridge University defines it well, stating, "As an AI system becomes more powerful and more general it might become superintelligent – superior to human performance in many or nearly all domains," ("Risks from Artificial Intelligence"). The word "general" here is in reference to one of two types of AI. The first type, narrow, describes AI that executes and masters one specific task, and it can perform this task better than any human could. However, narrow AI cannot do anything else except that assigned task. That is where the second type comes in, general AI. If an AI has reached general intelligence, which has yet to happen, it supposedly can perform any task at the same level as an average adult human. When AI *surpasses* human intelligence, it is said to be superintelligent, and that is where the problem lies.

The Singularity is predicted to bring about an unparalleled amount of change with it, and researchers cannot agree whether this change is good or bad. Some disagree on whether or not it may even occur. While such an event could wreak havoc on Earth, it could also be an opportunity to rid mankind of all strife. An intelligence of that level could potentially provide solutions that humans never would have thought of to major problems like cancer, world hunger, global warming and so on. The fear is that that same intelligence could also destroy us, because why would something smarter than us concern itself with our needs and values?

If a creation smarter than us chose not to support our existence, there would be little that could be done about it. If our goals do not align, what could we possibly do to challenge a being that considers us inferior? This is comparable to the way humans treat animals: we domesticate and breed animals the way we want to, hunt certain ones for food and sport, and leave others alone in the wild. Purely because we are smarter than them, we have power over them, and our actions have effects on their lives as well. Dodos are extinct because of over-hunting, rhinos endangered because we have killed thousands of them for their horns, and animals in zoos are kept under our thumb because we are more intelligent than they are. We have yet to experience the short end of the stick because we have always been the smarter species, but with AI, that may change.

The Singularity seems like an unlikely event, somewhere in the far off future, but Google's Director of Engineering and futurist, Ray Kurzweil, says otherwise. In a Futurism article, author Christianna Reedy reports that he has made a whopping 147 predictions from the 1990's, and of those predictions, eighty-six percent were correct. Regarding the Singularity, Kurzweil tells Futurism, "I have set the date 2045 for the 'Singularity' which is when we will multiply our effective intelligence a billion fold by merging with the intelligence we have created." The time he expects is a mere twenty-five years from now. Based on his accuracy, this may well become a reality.

Unfortunately, there is no known measure that exists today of when the Singularity will happen or how close we are to it, but an event of such magnitude demands our attention, time, and preparation. For a better picture, let us say that the history of our species is a lake, the future is the shore, and major events are stones; electricity would be a pebble, the invention of the

computer would be a palm-sized rock, and the introduction of a superintelligent AI would be a boulder — the waves will ripple out all across the lake, splashing the shoreline. As AI picks up momentum, we will start to see just how much weight it carries. Researchers and scientists from the beginning of AI development have speculated and projected many different ways this could go, and so AI's principal problems appear in their existing work.

A Brief History

To begin to understand AI, a basic history lesson is in order. An article titled "A Brief History of Artificial Intelligence," published in August last year, takes a look into the past. Authors Michael Haenlein and Andreas Kaplan, chronologically discuss the life of AI. The first functional electro-mechanical computer was created by scientist Alan Turing, and after that, research programs regarding AI were created in the 1950s (Haenlien and Kaplan). In the next decade, more working AI programs were formed, attempting to pass the Turing test of intelligence. Soon after this, the U.S. government took note and criticized high spending on AI research. Several years after that, deep neural networks were created artificially under the name of Deep Learning, and multiple tech giants are initiating similar projects to further AI research and applying that to their company's needs. The progress in AI in recent years has multiplied as computing power increases and information processing speeds up, which leads to predictions for the future. Many well-supported theories are being made — projections for what outcomes will follow AI's invasion of the world.

The Unemployment Problem

Possibly the most well-known projection of AI's effects is the forthcoming rise in unemployment. This issue is related more with AI's hardware counterparts, robots, that have yet to fully unite with smart software. Even so, the fear of joblessness is imminent. In the article, "Analysis of the Impact of Artificial Intelligence and Robotics on Human Labour" from the University of KwaZulu-Natal, the authors state, "It is apparent that artificial intelligence and robotics are substituting the previous generation's positions for fourth industrial [r]evolution jobs," (13877). Based on this, the replacement of humans in the workforce is not a matter of if, but when. The fourth Industrial Revolution is said to bring with it this massive disruption that will shake economies and governments.

Brian Alexander, an editor and writer of many technological texts, reports that Toledo, Ohio has more robots per worker than any other city in the United States. The fact that this is not a prediction of the future, but something that exists now is concerning to most. Alexander cites a study which states that "the state had lost 671,000 jobs to automation between 1967 and 2014, more than it lost to domestic competition and foreign trade combined" (22), which validates this concern among citizens. In Toledo alone, for every thousand workers, there were 2,374 robots in 2015. That was five years ago, meaning the number could have only grown now. This ratio is reflected across the nation as well, and in the article, "How Robots Are Beginning to Affect Workers and Their Wages," it mentions that, "Since 2009, the number of manufacturing robots has more than doubled—from 0.813 per thousand workers to 1.974 per thousand workers,"

which agrees with Alexander's findings. Machines can do jobs like jeep-assembly in the Toledo factory much better and faster than humans. Why would we *not* employ their assistance?

Well, it seems that the benefits of this advancement are reserved only for the heads of the industries that profit from rising unemployment levels. Robots imbued with advanced Machine Learning (ML) capabilities encroach on blue-collar jobs as they become more technologically competent. They perform the tasks of laborers with increased speed, precision, and efficiency. Robots do not require a lunch break, a healthcare system, or a salary, so it would be a smart decision on the company's part to start replacing its human employees. Rory Van Loo, a professor at Boston University School of Law, expands on this issue in his article, "Digital Market Perfection." The economy will be severely impacted by the rise of robots, crushing certain markets and allowing others to thrive. Loo says, "But in the extreme it could make some large markets more volatile, raising unemployment costs or financial stability concerns as more firms fail," and this inevitable disturbance in the workplace and the economy will create a crescendo in unemployment. Unfortunately, this is only one of the negative byproducts of AI's entry into society.

AI in Healthcare

AI has also entered society in the healthcare field, and at first glance, it seems like a major advantage. Patients may initially harbor some discomfort because the idea of an AI handling their private data, health care needs, and even their life is unnerving. AI's deep learning exists within a black box, meaning that the inner workings of each neural network — each code of the self-writing software — is not fully known, and that uncertainty is the root of this

discomfort. What we fail to realize is that the medicinal world is more or less a black box anyways to the general public, especially to individuals that are far from the medical field. A publication by the *Stanford Law & Policy Review* titled, "Artificial Intelligence in the Health Care Space: How We Can Trust What We Cannot Know," elaborates this notion. The authors state, "Health care always has been a place where individuals must put their faith in that which they do not fully understand," (Feldman, Aldana, and Stein), going on to describe that the nature of medicine itself is a black box. Despite this, the question still remains as to if patient discomfort is justified.

As an answer to that, A 2019 *Harvard Data Science Review* article by Michael Jordan specifies an instance in which those unknown networks of AI have proven to be harmful. Jordan provides a personal example in which an AI program mistakenly identified his unborn child as having Down Syndrome. The higher pixelation of the software's images revealed white spots near the fetus's heart, known as echogenic intracardiac focus (EIF), which is a common tell for Down Syndrome, but the imaging was faulty (Jordan). If Jordan had not figured that out with research of his own, his wife would have been subjected to an amniocentesis, which is a sampling of the fluid in the amniotic sac to spot abnormalities early on. This process comes with a bit of risk that could harm the mother, the child, or both, and it was only because of Jordan's precautionary measures that his wife did not go through with it. This misidentification is only one case. This one example sheds light on the many unknown mistakes that AI can make, because although machines tend to be more precise and accurate than humans, they are not perfect.

Discriminatory AI

AI's imperfection is expected since flawed humans are its point of reference. Our flaws, our society's flaws, namely discrimination, are adopted so quickly and extremely since AI learns so fast. In the article, "Let's Talk About Race: Identity, Chatbots, and AI," more of this adopted discrimination is discussed. Authors Ari Schlesinger, Kentan O'Hara, and Alex Taylor have contributed to the article from each of their respective specializations, with Schlesinger being from a tech institute, O'Hara conducting Microsoft Research that is relevant to the article, and Taylor bringing in the human-computer interaction aspect. One of the significant problems that they highlight with chatbots, or social communicating programs, is the issue with the data that it is fed. To be able to talk to humans, these chatbots must digest a large library of vocabulary words, phrases, regionalisms, curse words, idioms, and more, because human language is just that complex.

Chatbots have large databases from which they draw from to be able to talk and respond to users. In these databases, word filters are often placed to avoid offensive speech, and those filtered words are "blacklisted." However, the authors point out a problem with common worldfilters, saying, "In addition to the n-word, a blacklist may include strings like jap, paki, and homo; using these word-stems to catch hate-speech variants" (Schlesinger, O'Hara, and Taylor), which inhibits other words in the English language. The intention of these filters, of course, is to avoid discriminatory behavior, but without saying "paki," Pakistan cannot be spelled. Without saying "homo," words like homosexual, homogenous cannot be spelled either. If they were smaller, insignificant words, then maybe this could be overlooked but Pakistan is an entire

country and words like homosexual are used rather commonly. Blacklisting these prefixes has unfavorable consequences.

Not only this, but despite wordfilters being in place, there have been several cases of discriminatory AI. Tay, a Microsoft AI chatbot, became extremely racist, sexist, and anti-Semitic within twenty-four hours of being online. The cause of this is the information that it was fed — Internet sources, skewed data, social media, and discriminatory practices of humans. Racist tendencies, and discrimination in general, are practiced prominently on the Internet where identity can be hidden, making discrimination easier. The bots pick up on this and exhibit the same behavior. Chatbots also reflect the discomfort of many people when it comes to topics like race. It is noted in the article that an experimental chatbot avoided the topic when addressed and explicitly dodged any questions or conversations about race. When asked why, it seems to randomize an unrelated answer, much like a human child. The results discussed in this article confirm that AI could perform its tasks while harboring prejudice based on race, sex, and other various identities.

An example of explicit prejudice is described in Shelly Fan's book, *Will AI Replace Us?*, in which she mentions that "Google's first generation of its automated photo-labeling system sparked outrage when it misidentified people of African descent as gorillas," (Fan 76). Historically, the association of African people with gorillas comes from a discriminatory tactic employed by European conquerors that makes it seem like colored people are subhuman and, therefore, comparable to animals. The AI, incapable of distinguishing right from wrong, adopted this association and displayed it, misidentifying people with dark-colored skin.

Self-driving Cars

The misidentification does not stop there. Fan goes on to discuss a similar issue with self-driving cars. In 2016, 2017, and 2018, self-driving cars mistook objects in their surroundings as inanimate which led to accidents and loss of life (Fan 67). Without perfect programming that can adapt to any environment, it is apparent that AI can miscalculate its actions just like us. Fan clarifies that despite these collisions, software-driven cars still reduce casualties by a lot compared to human-driven cars, but with these recent accidents, public trust has tanked. People are easily frightened when some new volatile technology begins to err, and though statistically it is expected to significantly curtail road accidents, the fear has been established. The benefits of AI cannot be reaped because of the lack of trust, but too much trust can lead to reliance.

Smart Houses and Privacy

This dilemma is not limited to self-driving cars, but also a new technology that Patrick Lucas Austin introduces in his Time Magazine article, "Home is Where the Chip is." Currently, we have appliances such as smartphones, smartwatches, smart fridges, and automated assistants such as Amazon's Alexa and Google Home, but Austin debuts what will be known as smart-houses. These houses would be built from the ground up completely technologically compatible, with AI embedded in the walls and floors. In such a dwelling, inhabitants could step into the tub and the shower would know their preferred water temperature. When someone is showing symptoms of a cold or fever, the house AI would automatically order cold medicine beforehand. Such a world is a dream to some and a nightmare to others, which is understandable.

On one hand, it would be extremely convenient to have the house do everything for you, but on the other, human privacy is all but gone. Tasks like cooking, cleaning, and organizing would be assisted by robots, rendering the human unnecessary for these basic household activities. When human-computer interaction is no longer required, the network becomes an "Internet of Things" (IoT) system. Austin cites that consumers will be spending at least $123 billion in the IoT by 2021, and this will result in us slowly immersing ourselves into a world where nothing is left unautomated.

The risks and disadvantages that come with this future are primarily privacy concerns. With smart-houses could come smart-hackers that could infiltrate your house without a turn of the doorknob. Austin mentions that Daniel Cooley, chief strategy officer at Silicon Labs, says, "I call it a mass-extinction event for the Internet of Things," regarding the aftermath of a massive security breach in smart-houses, should they happen. This relates to public trust that is important for technology use and how it is affected when safety cannot be guaranteed or when it is impermanent. Self-driving cars that crashed during test drives and Amazon Echos that record people's private conversations have already shown what happens when AI is misused. Of course, there should be some level of fear of smart software, but if we cannot lend a bit of trust, it stunts progress.

Austin branches off of Cooley's comment to say, "Ultimately, experts say, people will come to see smart-home technology as essential as electricity, refrigeration or air-conditioning," and this is where issues are bound to occur. Trust is such a volatile thing; too little of it, and we are at an impasse, but too much of it, and we become puppets. There has to be a line between needing AI desperately to function in our everyday lives — so much so that we cannot live

without it — and using AI to our benefit while remaining in charge. This can very quickly go

south if we depend on our machines to take care of us, but the possibilities of AI's potential for a

better future cannot be ignored. I believe that when human trust is balanced, and AI does not get

out of hand, this issue can be averted. There is nothing wrong with wanting your shower to know

your preferred water temperature, but you should not need it like you need water to survive.

AI Trust Issues

This balance between trust and reliance is at the core of the issues with AI. Margit Sutrop

examines this more deeply in her article, "Should We Trust Artificial Intelligence?" Sutrop first

distinguishes between trust and reliance, iterating that trust is the cornerstone of AI. She states,

"In April 2019 the European Commission High Level Expert Group on AI adopted the

EthicsGuidelines for Trustworthy AI, stressing that human beings will only be able to

confidently and fully reap the benefits of AI if they can trust the technology." (499) She warns,

however, against letting that trust mutate into reliance because reliance creates room for

weakness and dependence, shifting the paradigm to us needing AI rather than AI needing us.

This same issue can lead to not only a lack of independence, but also a lack of original thought.

Ben Dickson features a college class's discourse about this in his article, "Why AI Leads us to

Think Less, Act Impulsively." Dickson reports that the professor of the class, Bernhardt Trout,

says that "these algorithms push us toward us thinking less and acting based on impressions that

may or may not be correct, as opposed to [making] our own decisions in a fully informed way. In

general, we want to have the answer and move on. And these algorithms tend to play off on that

psychology." Clearly, AI can take advantage of us if we relinquish free thought and take

everything for granted, and that is dangerous because it begins to scrape at our individuality and allows AI to determine how we live.

AI Arms Race

All of these issues are due to AI's relative innocence and its lack of an awareness of self, or a sense of right and wrong. AI is childlike, and yet it is being given far more power than it should be allowed. On the International scale, an AI arms race is underway, and China, Russia, and the U.S. are taking part. The article, "The Clash of AI Superpowers" calls attention to this issue that is being swept under the rug. In it, author Graham T. Allison states, "In fact, China stands *today* as a *full-spectrum peer competitor* of the United States in commercial and national security applications of AI," (11). This was in response to the common public opinion that no other country can outrank the U.S. in AI technology and militant capability, which is false. The Cold War threatened to end the world via nuclear bombs, but with the assistance of AI, it is much more feasible. Countries could penetrate each other's security measures digitally, and if the nuclear arsenal controls are within range, then off goes a nuclear war. The ignorance that government officials and certain scientists display to this globe-threatening arms race is concerning, and it is unclear whether this information is simply being hidden or if it is genuinely overlooked.

Jay Tuck, a defense expert, deliberates on the arms race specifically in the United States in his Ted Talk. He informs his audience that autonomous military AI is not simply in development, but many such technologies already exist and are in use. Atlas, a walking robot, plus his fellow siblings Big Dog, Wild Cat, Spot and more, all look like simple robots that are

learning to move across various terrains, designed by Boston Dynamics. However, Tuck notes

how suspicious it is that these machines are overseen by DARPA, the Defense Advanced

Research Projects Agency, and funded by Lockheed Martin, a powerful weapons company.

Surely, they are not solely interested in how these robots can walk, but rather their prospects in

replacing humans on the battlefield.

One of the more frightening creations shown in his speech is the Pegasus X47-B, a war

drone that the Navy owns. Tuck states, "It goes 2000 miles into enemy territory. It carries 2000

kilos worth of explosives, and it's run by Artificial Intelligence," (Tuck), all of which is

information that the general public was unaware of until recently. The Pegasus is run by AI that

can fly, administer an attack and return all on its own, and that is too much responsibility in the

hands of something that has yet to understand the concept of human life. Why is it that such

destructive power is given to a machine that lacks that fundamental notion to preserve and

protect? If, at some point, the driving program that runs the Pegasus were to misunderstand a

mission and attack the wrong people, or if it reached a point where it can decide for itself whom

to attack, the loss of life that will follow could be devastating.

Jay Tuck expands on such dangers in his article, "Mankind's Greatest Challenge:

Artificial Intelligence," in which he provides insight on what is considered the highest AI

advance of the decade - DeepMind. Google has inconspicuously been buying out robotics and AI

organizations, and DeepMind was one of them. Initially, it was supposed to be called Google

Brain, but they changed it to a less threatening name, DeepLearning. The DeepMind creators

were not originally Google folk, but they knew that they were in over their heads with the AI

tech that they were developing. These creators made Google agree to codes of ethics before

selling to them. The intricate neural networks of DeepMind have frightening prospects for the future regarding AI, because once the software starts writing itself faster and faster, at some point, we will no longer know what is being written.

Regulation

To prevent that from happening and to stymie AI from getting out of control, regulations have begun development, but they are in their preliminary stages — not yet universally applied. The Future of Life Institute is an organization that dedicates its research to major fields, namely AI, and is preparing rules by which programmers should handle AI. In their article, "Benefits and Risks of Artificial Intelligence," it states, "As these examples illustrate, the concern about advanced AI isn't malevolence but competence. A super-intelligent AI will be extremely good at accomplishing its goals, and if those goals aren't aligned with ours, we have a problem." ("Benefits and Risks…"). This is their justification for their concern and the measures they are taking towards it, and that misalignment of goals touches back to the most concerning result of AI — the Singularity. Unfortunately, not many organizations are following the Future of Life Institute, and they show little concern for AI's threatening presence. Our only hope is to create more regulations, improve our knowledge of deep AI networks, and find a way to keep our relationship with AI balanced.

Conclusion

AI is blazing a trail into our future, where the life we know today may no longer exist. Smart software is entering every aspect of human society, from massive tech giants, production

industries, and military institutions, to the very walls of the houses we live in. Self-driving cars, racist chatbots, smart-homes, autonomous weaponry — this is the future. Anything that is even remotely technologically advanced becomes something to question, something to fear, because our private lives are at risk of being exposed, and our civilizations are in danger of being transformed. These advancements are only the first steps of AI's coming ascent, and they confirm fears of the Singularity, which seems more and more inevitable as the days go by. The implications of every AI that exists to this day and those that will come into existence in the next few years are gargantuan. The ensuing chaos that we expect, are just the small issues that our human minds can think of, because we truly cannot predict the level of change that is to come.

The fear of our robot overlords is ever-present, and though it seems humorous, is it really that far fetched? The Singularity lurks in the horizon, taunting us with its mystery and shrouded with fear, as the ultimate destination of the road that AI leads us on. This game-changing creation that we have inflicted upon ourselves is certainly something to handle with the utmost caution, and although with all this negativity and dread it seems like it will be humanity's endgame, it does not have to be that way. Maybe a technological utopia is possible, with smart-houses that protect rather than expose our private data, with autonomous weaponry that have been taught to understand the values of humans and the worth of human life, with bots that recognize the history of discrimination that do not display negative behavior. It could just be the bridge to a new state of life.

Works Cited

Alexander, Brian. "The Politics of Angst in Robot City." *MIT Technology Review*, vol. 121, no.

4, pp. 20-25, Aug. 2018, *Academic Search Complete*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=a9h&AN=130151294&site=ehost-live&scope=site.

Allison, Graham T. "The Clash of AI Superpowers." *National Interest*, no. 165, pp. 11-24, Feb.

2020, *MasterFILE Premier*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=f5h&AN=141023237&site=ehost-live&scope=site.

Austin, Patrick Lucas. "Home Is Where The Chip Is." *Time Magazine*, vol. 194, no. 5, 5 Aug.

2019, pp. 42-47, *MAS Ultra - School Edition*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=ulh&AN=137701731&site=ehost-live&scope=site.

"Benefits and Risks of Artificial Intelligence." Future of Life Institute, 2016,

https:/futureoflife.org/background/benefits-risks-of-artificial-intelligence/?cn-reloaded=1.

Dickson, Ben. "Why AI Leads us to Think Less, Acti Impulsively." *PC Magazine,* pp. 27-30,

Jan. 2020, *MasterFILE Premier*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=f5h&AN=140975975&site=ehost-live&scope=site.

Fanafuthi, Mbali M., Jabulani Nyawo, and Pfano Mashau. "Analysis of the Impact of Artificial

Intelligence and Robotics on Human Labor." *Gender & Behavior*, vol. 17, no. 3, pp.

13877-13891, 2019, *Academic Search Complete*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=a9h&AN=139753131&site=ehost-live&scope=site.

Fan, Shelley. *Will AI Replace Us?: A Primer for the 21st Century.* Thames and Hudson, 2019.

Haenlein, Michael, and Andreas Kaplan. "A Brief History of Artificial Intelligence: On the Past,

Present, and Future of Artificial Intelligence." *California Management Review*, vol. 61,

no. 4, pp. 5-14, Aug. 2019, *Business Source Elite,* doi: 10.1177/0008125619864925.

Jordan, Michael I. "Artificial Intelligence - The Revolution Hasn't Happened Yet." *Harvard

Data Science Review*, 23 June 2019, doi: 10.1162/99608f92.f06Sch6oew61All.

Reedy, Christianna. "Kurzweil Claims That the Singularity Will Happen by 2045." *Futurism,* 5

Oct. 2017,

https://futurism.com/kurzweil-claims-that-the-singularity-will-happen-by-2045.

"Risks from Artificial Intelligence." *University of Cambridge*, 2019,

https://www.cser.ac.uk/research/risks-from-artificial-intelligence/.

Robin, Feldman, Ehrik Aldana, and Kara Stein. "Artificial Intelligence in the Health Care Space:

How we can Trust what we Cannot Know." *Stanford Law & Policy Review*, vol. 30, no.

2, pp. 339-419, 2019, *Academic Search Complete*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=a9h&AN=137238324&site=ehost-live&scope=site.

Rodgers, William M, and Richard Freeman. "How Robots Are Beginning to Affect Workers and

Their Wages." *The Century Foundation*, 17 Oct. 2019,

https://tcf.org/content/report/robots-beginning-affect-workers-wages/?session=1.

Schlesinger, Ari, Kentan O'Hara, and Alex Taylor. "Let's Talk About Race: Identity, Chatbots,

and AI." *City, University of London Institutional Repository*, 26 Apr. 2018,

http://openaccess.city.ac.uk/id/eprint/19124/1/Lets%20Talk%20About%20Race%20Acce

pted%20version.pdf.

Sutrop, Margit. "Should We Trust Artificial Intelligence?" *Trames: A Journal of the Humanities

and Social Sciences*, vol. 23, no. 4, pp. 499-522, 2019, *Academic Search Complete,* doi:

10.3176/tr.2019.4.07.

Tuck, Jay. "Artificial Intelligence: It Will Kill Us | Jay Tuck | TEDxHamburgSalon." *YouTube,*

uploaded by TEDx Talks, 31 Jan. 2017,

https://www.youtube.com/watch?v=BrNs0M77Pd4.

Tuck, Jay. "Mankind's Greatest Challenge: Artificial Intelligence." *Antioch Review*, vol. 77, no.

3, 2019, pp. 447-456, *MasterFILE Premier*, doi: 10.7723/antiochreview.77.3.0447.

Van Loo, Rory. "Digital Market Perfection." *Michigan Law Review,* vol. 177, no. 5, pp. 815-883,

Mar. 2019, *Business Source Elite*,

https://mulinutil1.millikin.edu:2443/login?url=http://search.ebscohost.com/login.aspx?dir

ect=true&db=bsh&AN=135585247&site=ehost-live&scope=site.