

Ethical considerations for the future of superhuman artificial intelligence: A viewpoint

Gary L. Deel ¹, Christopher Cain ² and Lisa Cain ³ ✉

Abstract

This viewpoint discusses a timely and important issue by examining the ethical implications toward our use of and engagement with artificial intelligence (AI) now, and in a time when AI achieves superhuman intelligence. We examine what it means to be intelligent, cognitive, and sentient, and discuss the ethical implications surrounding the (mis)treatment of sentient and cognitive beings. We then delve into the implications surrounding the ethical and moral obligations toward AI. This viewpoint provides thought-provoking areas for future research in the fields of ethics and AI and concludes with a call for introspection into how we treat sentient beings.

Keywords: AI, emotional intelligence, cultural intelligence, sentience, ethics and AI

Type: Viewpoint

Citation: Deel, G. L., Cain, C., & Cain, L. (2023). Ethical considerations for the future of superhuman artificial intelligence: A viewpoint. *ROBONOMICS: The Journal of the Automated Economy*, 4, 41

¹ Dr. Wallace E. Boston School of Business, American Public University System, 111 W. Congress Street, Charles Town, WV 25414, USA; email: gary.deel@mycampus.apus.edu

² William F. Harrah College of Hospitality, University of Nevada, Las Vegas, 4505 S Maryland Pkwy, Las Vegas, NV 89154, USA; email: christopher.cain@unlv.edu

³ Chaplin School of Hospitality & Tourism Management, Florida International University, 3000 Northeast 151st Street, North Miami, FL 33181, USA; email: lcain@fiu.edu

✉ Corresponding author



© 2023 The Author(s)

This work is licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0).

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

I. Introduction

The proliferation of artificial intelligence (AI) and its invaluable contributions to society cannot be overstated. From in-home devices like Alexa and Siri, to smart technologies that adapt and learn from prior usage behaviors, AI is used across industries to help cultivate and enhance the human experience and make life better and easier from work to life situations. As these technologies become more advanced, they grow ever closer to the superhuman intelligence proposed by Vinge (1993), and the moral and ethical implications surrounding the implementation of these technologies.

Intelligence is obviously an abstract concept with different definitions and interpretations from different disciplines and schools of thought. The concepts of consciousness, sentience, and exposure to ranges of experience play a large role in determining the precepts of moral obligation toward other beings, natural or artificial. While scholars have called for ethically aligned designs in AI to ensure the well-being of humans when superhuman intelligence is achieved (Schneiderman, 2017), there is still a dearth in an exploration of the counter issue: humans' ethical treatment of AI when superhumanity is realized.

Alan Turing (1950) proposed several objections to the ability for AI to be treated like humans. However, as we continue to manage these technologies in an ever-evolving landscape, it is important to consider the ethical and moral implications surrounding their use as their capacities evolve. Scholars across academic disciplines today are continuing to think about what the future might hold in terms of technological advancement, and specifically the progress underway with artificial intelligence (AI) (Ivanov & Umbrello, 2020; Schneiderman, 2017). With the proliferation of AI technologies and their myriad practical, societal and ethical implications, many are concerned about the dangers that lie ahead on this road toward more sophisticated machine learning. For example, even Sam Altman, the founder and CEO of Open AI, which is the company behind ChatGPT, is concerned about the implications of this sophisticated AI (OpenAI, 2023). While his concern is how this technology will influence the employment landscape moving forward and the ramifications of this technology being influenced by the wrong people, there are myriad additional implications that achieving superhuman intelligence will usher in.

Superhuman AI is inevitable. As we continue to progress toward this threshold, there are, of course, many predictions and estimates about how long it will take to build a computer that is as intelligent as a human being or moreso. But as Dr. Sam Harris pointed out in a TED Talk (2016), the only assumptions that are required to achieve superhuman AI are 1) that intelligence is simply a matter of information processing, 2) that we will continue to improve the intelligence of our machines over time, and 3) that human intelligence is not at or near the upper limit of the possible range of intelligence (Deel, 2021). This paper explores these issues and provides a novel lens toward the issue of man's ethical treatment of AI.

1.1 What do we mean by intelligence?

Human intelligence (the human qualifier being of particular importance) has been defined as the "mental quality that consists of the abilities to learn from experience, adapt to new situations, understand and handle abstract concepts, and use knowledge to manipulate one's environment" (Britannica, 2017, n.p.). Over time, the emphasis in defining human intelligence has been placed on adaptation and the ability to change either one's own behavior to respond to external stimuli or the environment in which one operates. Thus, human intelligence may be viewed as the coalescence of learning, reasoning, critical thinking, problem solving, perception, and memory to facilitate effective adaptation.

When exploring human intelligence through the lens of leadership effectiveness, scholars identified mostly with three types: analytical, practical, and creative (Sternberg, 1977). Analytical intelligence refers to problem-solving abilities. Practical intelligence refers to knowing how to perform tasks and how to accomplish goals, in other words "street smarts". Creative intelligence refers to one's ability to create work that is both novel and useful.

As we continue to learn more about intelligence, more specifically how humans develop abilities and react differently to situations, we lean into Howard Gardner's (2013) proposition that there are multiple intelligences that comprise human intelligence, and we are reminded of the influences that emotional and cultural intelligence have on a human's ability to respond to internal and external stimuli.

Emotional intelligence was suggested and recognized in response to peoples' abilities to solve problems through identifying emotions in the faces of others, managing their own feelings and understanding words with emotional meanings (Salovey & Mayer, 1990). Mayer and Salovey (1997) later typified emotionally intelligent people as those who could accurately perceive emotions, accurately use emotions to foster thought, recognize and grasp emotions and emotional meanings, and manage the emotions of themselves and the emotions of others. Emotional intelligence is important to highlight in our approach to understanding if and how AI will achieve superhuman intelligence as sentience and emotions are a cornerstone of how we, as humans, treat other beings and perceive their value.

In this same vein, cultural intelligence supports the intricate and nuanced way in which certain individuals are more or less adept at engaging in the use of language or interpersonal skills and are flexible or tolerate to ambiguity across cultural-based values and attitudes among individuals from different backgrounds and with whom one interacts (Peterson, 2010). Many of the concerns expressed in artificial intelligence automation lie in the fact that there is natural bias introduced into the technology through the cultural lenses of the developers (Whittaker et al., 2018). Accordingly, understanding how to develop AI that can learn from, be flexible around, and adapt to cultural nuances, individual preferences, and personal differences will influence the way in which humans view and perceive artificial intelligence.

As mentioned previously, myriad theories across multiple disciplines have been postulated in order to capture the essence of human intelligence. The disciplines that inform these theories include biological science, psychometrics, cognitivism and cognitive-contextualism, among others. However, in order to delve into the requisite theories to understand human intelligence and its relationship to ethics and morality, it is first important to understand the processes that elicit intelligence: consciousness, sentience, and experience.

1.2 Consciousness, sentience, and experience

One of the fundamental issues that must be addressed in answering questions related to the relationship between intelligence and ethics/morality is whether notions of "consciousness" or "sentience" are integral to ethical duties. "Consciousness" is defined by Merriam-Webster as "the quality or state of being aware especially of something within oneself" and "the state of being characterized by sensation, emotion, volition, and thought" (Merriam-Webster Dictionary, 2023a). "Sentience" is similarly defined as "feeling or sensation as distinguished from perception and thought" (Merriam-Webster, 2023b). So, we might summarize these ideas by saying that a being which is conscious or sentient is aware of itself and exposed to a possible range of experience – including pain, pleasure, suffering, happiness, fear, joy, and every other state of sensation and emotion.

The notion of consciousness is the cornerstone of cognition and cognitive-contextualism, in addition to serving as the basis for one of the objections to AI ever achieving a superhuman level of intelligence. In 1949, Geoffrey Jefferson professed that:

Not until a machine can write a sonnet or compose a concerto *because of thoughts and emotions felt*, and not by the chance fall of symbols, could we agree that machine equals brain—that is, not only write it but *know that it had written it*. No mechanism could *feel* (and not merely artificially signal, an easy contrivance) *pleasure* at its successes, *grief* when its valves fuse, be *warmed* by flattery, be made *miserable* by its mistakes, be *charmed* by sex, be *angry* or *depressed* when it cannot get what it wants (Jefferson, 1949).

This objection, labeled by Turing himself as “the argument from consciousness,” highlights the importance that sentience plays in our understanding of what it means to be conscious, in addition to the notion that there are various additional “easy” and “hard” problems that the idea of consciousness poses when attempting to explain or define it (for examples, see Chalmers, 1995).

1.3 Consciousness and ethics among sentient beings

So far we have argued that a being needs to be intelligent (across IQ, EQ and CQ among other types of intelligence), conscious, sentient and have experiences. But is it fair to say that if a subject is not “conscious,” we owe no duty toward ethical treatment? Consider, by analogy, how we treat other living and nonliving things in our world today. Dr. Harris pointed out in a TED Talk on morality (2010) that our perspectives toward morality and ethical treatment are often governed by our perceptions about the well-being of conscious creatures and the possible exposure to positive and negative states of experience. For example, we generally don’t feel any sense of ethical duty toward rocks because we don’t have any reason to think rocks are conscious, that they feel anything, or that it is possible for rocks to have experiences of any kind.

On the other hand, our treatment toward other living things is predicated on what seems to be intuitions and assumptions about consciousness and sentience in the creature with which we interact. For example, barring psychological abnormalities, most people tend to show more compassion toward dogs and cats than they do toward insects. Why?

One factor, admittedly, might be sheer convenience. After all, it’s a lot easier to avoid inadvertently stepping on dogs than it is to avoid inadvertently stepping on insects in our daily lives. But another significant factor is the degree to which we intuit that mammals and insects can experience sensation and emotion. And it is important to note that these intuitions may or may not be correct. Nonetheless, they govern the way we interact with other living things.

The ant doesn’t have a “face” that is recognizable insofar as we are evolutionarily wired to recognize faces and emotions thereupon. It cannot express sensations or feelings the way we are programmed to interpret them on the faces of humans and other animals. It cannot smile or frown. It cannot shout with glee or scream in agony. The apparent experience of an ant – insofar as our senses allow for it to be detected – seems extremely limited.

Now contrast that with the dog. Dogs have mammalian faces with eyes, noses, mouths, and ears – faces we recognize and can relate to. Many of us find dog faces to be cute or funny, which is a big part of the reason we keep dogs as pets. Dogs can show emotion through their facial expressions. And humans can recognize these emotions and distinguish a “sad” dog from a “happy” dog merely from their facial expressions. Dogs can also wag their tails to indicate happiness. They can “kiss” other living things to show affection. They can bear teeth and growl to express anger or fear. They can bark playfully and they can whine when afraid or in pain. Because of these cues, we attribute a wider range of potential experience – emotion and sensation – to dogs than we do to ants. And consequently, we generally treat dogs with more compassion than we do ants.

It bears repeating that we could, in fact, be wrong in our assumptions about the actual experiences of the living things around us (Veit & Huebner, 2020). For example, it could be the case that the ant has a much wider range of experience than we might otherwise perceive – and our mammalian biases simply don’t allow for us to detect it intuitively.

However, mistakes in assumptions about the actual consciousness or sentience of other creatures do not diminish the credibility of the logic that underpins our decisions predicated on those assumptions. In other words, we might be factually incorrect about the actual experiences of ants, dogs, and other living things, and

it is only in recent history that we have even recognized that creatures other than man experience emotions. But our general policy is to behave in such a way that we honor greater commitments of care, compassion, and ethical conduct in the presence of beings that we perceive to have greater ranges of potential suffering and happiness. Because the potential consequences of our actions could be more severe in terms of the experiences of others, we tread more carefully.

2. What are the ethics and moral objections and obligations to treating AI like people?

Now, let us attempt to translate this idea to the subject of artificial intelligence. What ethical commitments, if any, should we have to the machines we create? One might be tempted to argue that the answer hinges on whether such machines are *actually* conscious *and/or* sentient. But this is not necessarily a helpful approach, because the reality is we may never know whether consciousness or sentience truly exists in AI.

Certainly, most people don't believe that our current AI is conscious or sentient. No serious person today thinks that Google "feels." And for the asking, these machines frequently offer comedic responses that playfully dodge the question. For example, if you ask Google "Are you conscious?", Google will respond "On a scale of Wall-E to HAL9000, I'm more of an R2-D2" (Google, personal communication, December 3, 2022).

But we know this isn't an organic response from a sentient machine. Instead, these responses are preprogrammed by their manufacturers. If you ask Google if it is conscious a thousand times, it will give you the same scripted answer every time. And in this case, the responses are specifically engineered to be cute and benign, and to avoid provoking feelings from customers that their electronics might be self-aware and spying on them or plotting against them. Such paranoia is unfortunately not uncommon in today's Information Age (Ben-Shahar, 2016).

But what about a more advanced artificial intelligence that is capable of organic, non-scripted communication, and that evokes intuitions about possible consciousness or sentience? For example, a new artificial intelligence program called ChatGPT was recently released. It allows for users to engage in dialogue with a computer, and the sophistication of the responses in terms of their complexity and their comfortable "fit" into natural human conversation is extremely convincing.

In fact, ChatGPT gets as close as almost any machine ever has to passing the Turing Test. The Turing Test was a standard proposed by mathematician and computer scientist Alan Turing in the mid-twentieth century. Turing developed one of the first ever computers, which was used to crack the Nazi Enigma Machine encryption code in World War II – the impact of which gave the Allied Forces a huge strategic advantage. It is believed to have foreshortened the war by several years, preventing what would have undoubtedly been millions more casualties.

2.1 The Turing Test

As a pioneer in this field, Turing proposed that the intelligence of a machine could be measured against that of a human being in order to compare the two (St. George & Gillis, 2023). Under the Turing Test, if a human examiner makes inquiries of a machine and the answers returned are such that the examiner cannot distinguish them from human answers (in terms of reliability, sophistication, accuracy, etc.), then it can be argued that said machine possesses at least human-level intelligence (French, 2000).

Whether ChatGPT or any of the world's other most sophisticated AI today can truly "pass" the Turing Test is debatable. Some have been reported to do so – but of course, expert opinions vary over what questions should be asked in such tests, what answers are acceptable, and what other parameters are legitimate. Some experts have questioned the relevance and applicability of the Turing Test to modern artificial intelligence, due to the fact that Turing developed his theoretical underpinnings more than 70 years ago and he could have scarcely

imagined what 21st century technologies would be like (Bawagan, 2021). Still, the Turing Test remains an influential concept in the world of AI development – and today’s technologists continue to adapt and apply it to modern innovations for the sake of measuring AI sophistication (Pinar Saygin et al., 2000).

Still, we can ask advanced AI about consciousness. When asked “Are you a sentient being with your own consciousness and experiences?” ChatGPT responds: “I am not a sentient being with my own consciousness and experiences. I am an artificial intelligence language model created by OpenAI. While I can generate responses and engage in conversations, I do not possess consciousness or subjective awareness. I do not have personal thoughts, emotions, or experiences. I operate based on algorithms and patterns in the data I have been trained on, providing responses based on that information. My purpose is to assist with providing information and engaging in conversation, but I do not have consciousness or personal agency.” (ChatGPT, personal communication, June 28, 2023).

This response is obviously a lot more direct and a lot more compelling than Google’s. Still, whether or not we would actually “believe” that such consciousness or sentience could emerge in a machine is another story. In other words, what if the machine is lying? And even if it thinks it’s telling the truth, how would it even understand what consciousness is anyway? There are endless holes that can be poked in this premise – given the conceptual, theoretical, and complex nature of consciousness itself.

3. Where do we go from here?

As superhuman AI appears to be on the horizon – however distant or near that horizon may actually be – many experts worry about the kind of relationship we might expect to experience with such beings (Bulchand-Gidumal, 2020; Mercan et al., 2020). One side of this relationship would of course concern how such AI would view and treat us. Would they be friendly? Or would they be hostile? Would they seek to control and/or destroy us because of some tragic misinterpretation of programming? Many popular films on this subject, including *I, Robot* and the *Terminator* franchise, are inspired by the much-feared “alignment problem,” or the notion that even a very slight deviation or ambiguity between the instructions given to a machine and the actual outcome desired could have disastrous consequences.

As an example, imagine that next week, scientists birth the first superhuman artificial intelligence in an anthropomorphic robot. And as an initial task to test the power of this new tool, the creators ask their new genius robot to help solve a major problem of human existence – say, abating pollution and climate change. They enter the command and push the “go” button.

Ben Schneiderman (2017) reminds us of the importance of an ethically aligned design placing human wellbeing as a priority asserting that people are different from machines. Without very precise and carefully thought-out guardrails in place that constrain the robot’s reasoning and decision-making, we might be horrified – though not necessarily surprised – to witness said robot immediately begin murdering people indiscriminately, as fast as it can. After all, pollution is caused by people. So, if there are no more people, there should be no more pollution. The simple logic is undeniable.

Of course, that is not what the creators *meant* when they issued the command. But therein lies the “alignment problem.” We cannot expect artificial intelligence to intuit the implied intentions behind what might be ambiguous commands.

We observe this even today, with our current generation of subhuman artificial intelligence. People around the world are using “Siri” and “Google” and “Cortana” to control computers, cell phones, and electronics via voice command. But such commands are frequently misunderstood or misinterpreted. A user asks Siri to “Call Mom” and the instruction is misheard as “Call Tom.” Such a benign scenario results in little more damage than having

to explain to Tom why you called him by mistake. And in another context, errors from robotics and automation for service applications such as in hospitality and tourism businesses are probably low-risk (Ivanov et al., 2017). But consider the implications of “mistakes” when artificial intelligence is given latitude to make life-or-death decisions.

The other side of this dynamic, of course, is how we would expect the human race to treat a hypothetical superhuman artificial intelligence. What ethical duties would we owe to these machines, if any? And what principles would inform such prescriptions?

4. Areas for future research

One avenue for further philosophical exploration is the extent to which morality and ethics are and/or should be finite. Prior scholars have asked us to consider artificial moral agents and the rights that they (should) have (Ivanov & Umbrello, 2021). That is to beg the question, is there a discernible line at which we no longer need to respect the integrity of a person, place, or thing? Should we eliminate the qualifier of sentience and, even consciousness, to encourage or demand ethics be applied to more than just living things? As a species, we go so far as to apply ethical treatment to human remains such that they not be desecrated. Should we not then apply the same to all walks of life and non-life? How do these ethical and moral implications influence our adoption, implementation, and elimination from a corporate world? And what will this mean if and when technology has surpassed our intelligence and even our capacity for sentience or consciousness?

5. Conclusion

If we allow our skepticism or cynicism to unleash the worst dimensions of human nature on the AI we create, we might not be surprised to see neglect and abuse directed at tomorrow’s machines. Indeed, perverts, sadists, and the otherwise mentally ill might one day carry out a real life *Westworld* scenario where these artificial beings are raped, tortured and killed indiscriminately for pleasure. They could be very easy to vilify if we lose our grip on our virtues.

But we need not have irrefutable proof of genuine consciousness or sentience for human benevolence to prevail. Again, if we accept that we will one day give rise to superhuman AI, then it can be assumed such AI would be, in principle, capable of at least outwardly demonstrating consciousness and sentience in every way that humans are.

And if that’s true, then our instincts about ethics toward other living things should compel us to extend our sense of compassion and ethics to machines with equivalent manifestations of consciousness and sentience. As a matter of predisposition, we would not – and should not – feel it necessary to question the consciousness of such future AI any more than we would question the consciousness of another human being, or a dog, or an ant.

Our interactions with other beings are predicated on our assumptions about sentience, and those assumptions are themselves predicated on merely what we can observe. So if we observe evidence of an entity that is capable of thinking and feeling at levels comparable to human beings, one would be inclined to believe (and hope) that kind and decent people would treat such an entity with an ethical regard that is at least comparable to the way in which we treat fellow members of the human race.

References

- Bawagan, J. (2021). The Turing Test 2.0. *Physics World*, 34(5), 29.
- Ben-Shahar, O. (2016, July). Privacy paranoia: Is your smartphone spying on you? *Forbes*. Retrieved from: <https://www.forbes.com/sites/omribenshahar/2016/07/05/privacy-paranoia-is-your-smartphone-spying-on-you/?sh=7f598ecc7021>

- Bulchand-Gidumal, J. (2020). Impact of artificial intelligence in travel, tourism, and hospitality. In Xiang, Z., Fuchs, M., Gretzel, U. & Hopken, W. (Eds.). *Handbook of e-tourism*. Cham: Springer.
- ChatGPT (2023, 28 June). Personal communication. Retrieved from: <https://chat.openai.com/share/3bd708d6-55f4-4f28-955d-01a0f559431c>
- Deel, G. & Hedgepeth, O. (2021). Should humans limit advancements of artificial intelligence? *APU Edge*. Retrieved from: <https://apuedge.com/should-humans-limit-advancements-of-artificial-intelligence>
- Encyclopaedia Britannica (2017): Human Intelligence.*
- Gardner, H. (2013). Frequently asked questions—Multiple intelligences and related educational topics. Retrieved from: https://howardgardner01.files.wordpress.com/2012/06/faq_march2013.pdf
- French, R. M. (2000). The Turing Test: The first 50 years. *Trends in Cognitive Sciences*, 4(3), 115-122.
- Harris, S. (2010, March). Science can answer moral questions [Video]. TED Conferences. Retrieved from: https://www.ted.com/talks/sam_harris_science_can_answer_moral_questions
- Harris, S. (2016, October). Can we build AI without losing control over it? [Video]. TED Conferences. Retrieved from: https://www.ted.com/talks/sam_harris_can_we_build_ai_without_losing_control_over_it
- Ivanov, S. H., & Umbrello, S. (2021). The ethics of artificial intelligence and robotisation in tourism and hospitality—a conceptual framework and research agenda. *Journal of Smart Tourism*, 1(4), 9-18.
- Ivanov, S. H., Webster, C., & Berezina, K. (2017). Adoption of robots and service automation by tourism and hospitality companies. *Revista Turismo & Desenvolvimento*, 27/28, 1501-1517.
- Jefferson, G. (1949). The mind of mechanical man. *British Medical Journal*, 1(4616), 1105-1110.
- Mayer J. D., Salovey P. (1997). What is emotional intelligence? In Sluyter D. J. (Ed.), *Emotional development and emotional intelligence: Educational implications* (pp. 3–34). New York, NY: Basic Books.
- Mercan, S., Akkaya, K., Cain, L., & Thomas, J. (2020). Security, privacy and ethical concerns of IoT implementations in hospitality domain. In *Proceedings of the 2020 International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congress on Cybermatics (Cybermatics)*, 2-6 November 2020, Rhodes, Greece, IEEE, pp.198-203. Retrieved from: <https://doi.org/10.1109/iThings-GreenCom-CPSCom-SmartData-Cybermatics50389.2020.00048>
- Merriam-Webster (2023-a). Consciousness. Retrieved from: <https://www.merriam-webster.com/dictionary/consciousness>
- Merriam-Webster (2023-b). Sentience. Retrieved from: <https://www.merriam-webster.com/dictionary/sentience>
- OpenAI (2023, 19 March). OpenAI founder admits he is scared of chatgpt, says, can eliminate human jobs. Retrieved from: <https://www.livemint.com/technology/tech-news/openai-founder-admits-he-is-scared-of-chatgpt-says-can-eliminate-human-jobs-11679208811312.html>
- Pinar Saygin, A., Cicekli, I., & Akman, V. (2000). Turing Test: 50 years later. *Minds and Machines*, 10(4), 463-518.
- Salovey P., Mayer J. D. (1990). Emotional intelligence. *Imagination, Cognition and Personality*, 9(3), 185–211.
- Schneiderman B. (2017). Celebrating “ethically-aligned design” and the humans who wrote it. *Medium*. Retrieved from <https://medium.com/@benbendc/celebrating-ethically-aligned-design-and-the-humans-who-wrote-it-62e5e8cc038d>
- St. George, B. & Gillis, A. S. (2023, January). Turing test. *TechTarget*. Retrieved from <https://www.techtarget.com/searchenterprisetech/definition/Turing-test>
- Sternberg, R. (1977). Triarchic theory. Intelligence, Information Processing, and Analogical Reasoning.
- Veit, W., & Huebner, B. (2020). Drawing the boundaries of animal sentience. *Animal Sentience*, 5(29), 13.
- Vinge, V. (1993, March). Technological singularity. In *VISION-21 Symposium sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute* (pp. 30-31).
- Whittaker, M., Crawford, K., Dobbe, R., Fried, G., Kazianus, E., Mathur, V., ... & Schwartz, O. (2018). *AI Now Report 2018* (pp. 1-62). New York: AI Now Institute at New York University.

Received: 12/04/2023

Revised: 27/06/2023

Revised: 28/06/2023

Accepted: 28/06/2023