International Journal of Scientific & Technical Development
(A Peer-Reviewed Bi-Annual Journal) ISSN: 2348-4047

ETHICS AND ENGINEERING: A STUDY OF ARTIFICIAL INTELLIGENCE

***Dr. Gurinder Kaur Sodhi, **Ifat Rasheed**
**HOD, Department of ECE, Desh Bhagat Univesity, Punjab.
**Research Scholar, Department of ECE, Desh Bhagat Univesity, Punjab.

.

Abstract
The hot topics among academics and IT researchers include artificial intelligence and challenges linked to it. The goal of the researchers is to gather and repurpose the previously existing, constantly expanding information via various social media platforms. The abundance of data has increased study into its processing, analysis, and application. The emphasis is on using computers to handle data more effectively because they are far faster and better at it than people. This increasing enthusiasm for research, academia, business, and open source has produced innovations and advancements that have the potential to bring about significant change. However, the prospect of building intelligent robots presents a few moral questions regarding the machines' own moral standing as well as how to make sure they don't damage people or other morally significant entities. As a result, in contemporary culture, ethics and law are closely connected, and many legal judgments are based on how various ethical concerns are interpreted. These problems take on a new dimension thanks to artificial intelligence. Artificial intelligence-based systems are evolving toward more autonomy in terms of the complexity of the tasks they can do, the potential influence they may have on the world, and the decreasing capacity of humans to comprehend, forecast, and manage their operation. Most people don't realize how automated these systems really are; they may learn from their own mistakes and take actions that go beyond what their designers had in mind. There are several moral and legal issues with this. A feeble attempt has been made to critically assess the ethical concerns surrounding artificial intelligence and thinking machines throughout the world in this article.
Key Words: Artificial Intelligence; Thinking Machines; Ethical Issues; Science and Technology.

Introduction:

Although AI algorithms have beaten humans in many specific domains such as chess, there is nearly universal agreement among modern AI professionals that AI falls short of human capabilities in some critical sense. It has been suggested by some that as soon as AI researchers figure out how to do something, that capability ceases to be regarded as intelligent. It is debatable whether human intelligence is truly general—we are certainly better at some things. When human engineers construct a nuclear reactor, they imagine the exact events that may occur inside of it (such as valves failing, computers malfunctioning, or cores becoming hotter) and design the reactor to make these occurrences unlikely to be disastrous. The toaster-paradigm, the realm of locally preprogrammed, particularly envisioned behavior, is broken down even by task-specific AI algorithms. The programmers would have had to manually preprogram a database containing moves for every potential chess position that one may face since the machines can only behave exactly what they are instructed. To serve the ultimate goal of feeding oneself, modern people actually perform millions of things. Few of these pursuits were "envisioned by Nature" in the sense that they represented challenges from our ancestors to which we were specifically equipped. However, our evolved brain has become sufficiently potent to be substantially more widely applicable; to allow us to predict the results of millions of varied acts across domains and to influence the ultimate outcomes. Despite the fact that none of our predecessors faced a problem comparable to vacuum, humans have traveled across space and left footprints on the Moon. Designing a system that would function securely in thousands of situations, including circumstances not precisely envisioned by either the creators or the users, including contexts that no human has yet encountered, is a qualitatively different task from domain-specific AI. There may not be a compact local description of all the ways that people get their daily bread in this situation, nor is there a local definition of good behavior—a simple specification over the actions themselves. One must describe appropriate behavior in such a manner that they may not hurt humans in order to create an AI that behaves safely while functioning in various domains, with numerous repercussions, including those the engineers never expressly envisioned. This is projecting the far-reaching effects of activities, making it non-local. As a result, this specification can only be achieved as a design attribute and be successful if the system explicitly extrapolates the effects of its behavior. From the standpoint of public relations, this may not seem like a desirable situation, but it's diffi-

cult to see how a general intelligence functioning on unanticipated issues across domains with preferences over distant repercussions could be guaranteed to behave ethically in any other way. However, we couldn't foresee which precise answer the mind would find. By looking at the cognitive design, we might be able to confirm that the mind was actually looking for alternatives that we would categorize as ethical. Respecting such a verification necessitates a means of separating sincere guarantees from wishful thinking and irrational optimism (a technique that won't claim the AI is safe until the AI is truly safe). Keep in mind that artificial intelligence research has historically struggled with unrealistic expectations. It will need an AGI that thinks like a human engineer concerned with ethics, not simply a basic product of ethical engineering, to develop a trustworthy AGI that can be verified. This will require new techniques and a different style of thinking from checking power plant software for faults. As a result, the field of AI ethics, particularly as it relates to AGI, is likely to be fundamentally different from the field of noncognitive technology ethics in the following ways:

1. Even if the programmers follow all the rules, the local, specific behavior of the AI might not be predictable apart from its safety.
2. Verifying the safety of the system becomes more difficult because we must verify what the system is trying to do rather than being able to verify the system's safe behavior in all operating contexts.
3. Ethical cognition itself must be taken as a subject matter of engineering.

Superintelligence can be developed to be useful, and depending on its technical prowess, it may be able to quickly solve a number of current issues that have eluded our human-level intelligence. One of the several "existential risks" listed by Bostrom (2002) is superintelligence: a risk "where an adverse outcome would either annihilate Earth-originating intelligent life or permanently and drastically curtail its potential." On the other hand, a successful conclusion for superintelligence may protect and advance intelligent life that originated on Earth. It is crucial to stress that there are significant potential hazards and advantages associated with having better minds. Various cognitive biases, such as the "good-story bias" put out by Bostrom (2002), may make it difficult to reason about the likelihood of global catastrophes (Yudkowsky 2008b). Consider the possibility that our perceptions of what future situations are "plausible and realistic" are influenced by the world around us.

Although Asimov's Three Laws of Robotics (Asimov 1942) are occasionally used as a guide for developing moral AI, they serve as little more than a story device than Asimov's "positronic brain." Asimov wouldn't have had any tales if he had shown the Three Laws as successfully operating. Thinking of "AIs" as a species with set traits and wondering if they would be nice or evil is a mistake. Since all people share a similar brain architecture, the phrase "Artificial Intelligence" alludes to a broad design area that is probably far greater than the space of human brains. Asking, "Will AIs be good or evil?" as though attempting to choose a concept for a movie narrative may be a sort of good-story bias. "Exactly which AI design are you talking about?" should be the response. Can one have influence over an artificial intelligence's early programming to affect how the world will be affected later?

How can you create an AI that, when it operates, becomes more ethical than you is the ultimate machine ethics dilemma. It is no longer comparable to asking our own philosophers to develop superethics. However, we must be able to clearly state the query if not the solution. Machine ethics must dedicate itself to attaining human-superior (not simply human-equivalent) niceness if robots are to be put in a position of becoming stronger, quicker, more trustworthy, or smarter than humans.

Legal constraints involved in AI across the globe:

In the case of robots in particular, the legal issues are significantly more severe. Predictability is essential to current legal techniques, and a system that learns from information it gets from the outside world may behave in ways that its developers could not have foreseen. Furthermore, such systems may run independently of their designers or operators, making it more difficult to assign blame. These traits raise issues with predictability and the capacity for independent action while absolving oneself of accountability.

Several possibilities exist for regulation, including regulation based on already-established norms and standards. Artificial intelligence-based technology, for instance, may be governed as either property or as goods subject to copyright. However, issues occur when we consider that such technologies have the capacity to behave independently of their designers, owners, or proprietors. Since animals are likewise capable of autonomous action, it is conceivable to apply the laws that govern a certain type of ownership in this regard.

Although they are somewhat constrained, proposals on

how the legislation on animals should be applied have been offered. First off, it is improper to apply laws based on analogy inside the confines of criminal law. Second, it is reasonable to assume that domestic dogs, for whom these regulations were initially established, will not injure people ordinarily. Since the laws controlling the keeping of wild animals are more strict, there have been requests for comparable laws to be applied in more advanced legal systems. However, the issue here is how to distinguish between the many aspects of artificial intelligence that were previously described. Furthermore, strict regulations may potentially delay the adoption of artificial intelligence technology because of the unanticipated dangers of creator and inventor liability.

Another common approach is to use the same standards to govern how legal companies conduct their business. Robots can be accorded a legal status equivalent to that of a legal entity since they are artificially created subjects of the law. Anybody can be given rights if the law is sufficiently forgiving. Additionally, it may limit rights. For instance, historically speaking, slaves essentially had no rights and were seen as property. It is also possible to witness the converse scenario, in which rights are granted to things even when they don't explicitly show that they can do anything. Even today, both in wealthy and developing nations, there are instances of strange items that are acknowledged as legal entities. In a similar vein, robots can be held accountable for their activities without requiring them to have free will or intentionality.

The analogy of legal entities, however, presents difficulties since swift and efficient administration of justice depends on the existence of legal entities. But despite the fact that it is hard to pinpoint who they are, the activities of legal entities may always be traced back to those of a single individual or group of people. In other words, businesses and other comparable entities are accountable for the deeds committed by their agents or workers. The actions of artificial intelligence-based systems will not necessarily be traced back to the actions of an individual, and legal norms on the sources of increased danger can be applied to artificial intelligence-based systems. The challenge is identifying which artificial intelligence systems can be deemed criminally responsible.

The formulation and promotion of a strategy for the sustainable growth of the smart robot industry is how the law hopes to improve living conditions and advance the economy. The government develops a fundamental strategy to guarantee the accomplishment of these objectives every five years. Similar to this, in 2018 President Emmanuel Macron unveiled France's new national artificial intelligence plan, which calls for spending 1.5 billion Euros over the following five years to assist the field's research and innovation. The plan is based on the suggestions included in the report written under the direction of French mathematician and National Assembly delegate Cédric Villani. The decision was taken to focus the approach on four distinct industries: security, transportation, healthcare, and the environment and environmental protection. The justification for this is to concentrate the potential of artificial intelligence's comparative advantages and competences on industries where businesses may play a significant global role, as well as because these technologies are crucial for the general good.

Seven important recommendations are made, one of which is particularly pertinent to the goals of this article: making artificial intelligence more approachable. It's accurate to say that the algorithms utilized in artificial intelligence are specific and, frequently, trade secrets. Algorithms can, however, be biased. For instance, during the self-learning process, they may absorb and accept societal prejudices or those that are imparted to them by developers and base their judgments on them. This has been done before in the law. Based on data from an algorithm that forecasted the risk of repeat offenses, a defendant in the United States was given a hefty jail term. The parameters used to assess the likelihood of repeat offenses were a trade secret, thus they were not revealed in the defendant's appeal against the employment of an algorithm in the sentencing procedure. The French approach suggests establishing an ethics advisory group, specifying the ethical duty of people involved in artificial intelligence research, and constructing transparent algorithms that can be evaluated and confirmed.

The situation with the European Union is the same. The first step toward regulating artificial intelligence in the European Union was the formulation of the resolution on the Civil Law Rules on Robotics. In 2015, a working group was created in the European Union to address legal issues relating to the advancement of robots and artificial intelligence. The resolution does not have legal force, but it does make certain recommendations to the European Commission about potential measures in the field of artificial intelligence, both in terms of civil law and in terms of the moral implications of robotics.

A "smart robot" is one that has autonomy through the use of sensors and/or connectivity with the environ-

ment, has at least a minimal amount of physical support, adapts its behavior and actions to the environment, and cannot be said to have "life" in the biological sense. It is suggested to "introduce a system for registering advanced robots that would be managed by an EU Agency for Robotics and Artificial Intelligence." Both strict responsibility (no fault necessary) and risk-management approaches (liability of a person who was competent to minimize the risks) are presented as solutions for culpability for damage brought on by robots. The resolution states that liability "shall be commensurate to the actual level of instructions issued to the robot and to its level of autonomy. A mandatory insurance program for robot users and a compensation fund to pay out compensation in the event that no insurance policy covered the risk might supplement liability rules. A Code of Ethical behavior for Robotics Engineers and a Code for Research Ethics Committees are proposed in the resolution as two standards of behavior for handling moral dilemmas. The first code suggests the following four standards for robotics engineering ethics: Robots should be beneficent (act in humans' best interests), non-malevolent (do no damage to people), autonomous (human connection with robots should be voluntary), and just (fairly divide the advantages of robotics).

Artificial intelligence (AI) and the technology that uses it, machine learning (ML), provide the biggest ethical problem. The Hanson Robotics robot Sophia has helped bring this topic much more into the public's notice. Sophia was recently awarded citizenship in Saudi Arabia, a development that pushed AI and ML and the difficult ethical issues they raise into the public eye. Sophia became well-known and was invited to appear on discussion programs all around the world. She is frequently asked if she believes that robots will wipe out mankind. The ML applications that are less obvious are more significant, though. Sophia and other glorified chatbox robots like her are not going to end the world, but we should be worried about the plethora of new, tough ethical questions that AI and ML are now posing globally. States will be more eager to invest in other facets of this technology, which will unfortunately be riskier and more morally dubious. States are primarily interested in applying AI and ML in three areas: military, intelligence, and judicial systems (including law enforcement and court systems). The machines create their own models on which to act or judge by definition. Because governments will only be able to defend judgments made using these algorithms' advice to a limited extent, transparency is a critical challenge. This may make transparency less significant in democracies. This may allow people in charge to disregard requests for openness, which might be problematic for reformers or democratic campaigners. In countries like China, where transparency isn't even acknowledged, AI is already being used in law enforcement. Artificial intelligence is used with facial recognition technology to more precisely measure and define face traits, enabling security cameras to recognize jaywalkers. AI will already be used in China to foresee terrorism and societal upheaval. However, these phrases have extremely varied meanings in other nations. Even if certain definitions of terrorism and civil unrest may be morally dubious with or without AI and ML, these technologies will help states act on these definitions more effectively regardless of how morally dubious they may be. Courts in the United States are increasingly relying on risk assessment algorithms to identify criminal danger. Because there is no national or international regulation of this type of technology, countries will continue to equip their police forces and criminal justice systems with AI and ML-powered technologies without adequately addressing ethical concerns.

The ethical ramifications of AI and ML technology will likely only be discussed once it has been militarized and weaponized, similar to the scenario with nuclear power. The attraction is clear given that fewer soldiers are needed (resulting in fewer losses), efficiency increases (lowering long-term costs), and conflicts are easier to win. Although other countries, like as the United States, currently prohibit completely autonomous weapons, the risk still exists. Artificial intelligence is the future, according to Russian President Vladimir Putin. Whoever assumes control of this arena will also assume control of the whole planet. Russia has already made investments in the creation of a missile that is AI-powered. The concern is that a country would let a weapon like this to both identify something as a target and to fire at it without requiring human consent, even though it is unlikely to happen anytime soon.

The introduction of ethical standards along with global interventions into the current research and investment landscape are required to counteract and thwart the escalating arms race between the US, China, and Russia if we are to overcome the difficulties posed by machine learning technology. The best chance we currently have of controlling this issue before it permeates nearly every aspect of our lives, short of temporarily halting all public-private investments in and contracts utilizing machine learning technologies, is to create an

international organization akin to the IAEA to develop recommendations for applying human rights laws to machine learning technology, monitor nations infringing on human rights with machine learning technology, a. Instead of taking financing from institutions like the CIA for its research, the creation of an organization with a sizable budget devoted to machine learning technology would encourage private sector activities towards constructive uses. If states agree to only conduct machine learning technology research that has been deemed peaceful by this organization and to abide by certain civic obligations, it may eventually lead to a treaty similar to the Nuclear Non-Proliferation Treaty, which would prevent the kind of AI theater that Saudi Arabia has already engaged in. In the past, it has taken a crisis or large-scale catastrophe to establish these types of agreements and organizations. We have the chance for prevention using ML. Given how much they stand to gain from a world in which AI is unregulated, persuading the U.S., Russia, and China to support these solutions would be another issue, but it is obvious that the implications will be severe if they don't.

Ethics and Artificial Intelligence

The moral obligations we have in our interactions with modern AI systems are all based on our duties to other creatures, such our fellow humans, and not on any obligations to the systems themselves. While it is generally agreed that modern AI systems lack moral status, it is not obvious exactly what characteristics moral status is based on. Sentience and sapience (or personhood) are two qualities that are frequently put up as being crucially connected to moral standing, either independently or in combination. These can be loosely described as follows:

Sentience is the ability to have remarkable experiences or qualities, such as the ability to experience pain and suffering. Sapience is a group of abilities linked to higher intellect, including self-awareness and the ability to reason- a flexible agent One widely held belief is that while many animals possess qualia and hence some moral standing, only humans possess sapience, elevating them to a higher moral standing than other non-human creatures. This perspective must, of course, deal with the existence of borderline cases, such as, on the one hand, infants or people with severe mental retardation, who are sometimes regrettably referred to as "marginal humans," who do not meet the requirements for sapience, and, on the other hand, some non-human animals, like great apes, who may have at least some of the traits of sapience. People who label themselves "marginal humans" contest their moral position. Others have suggested additional criteria that an item may meet in order to qualify as a bearer of moral status, such as belonging to a species that typically possesses sentience or sapience or being in a proper relationship to a creature that has moral status on its own (cf. Warren 1997). But for now, let's concentrate on the sentience and sapience standards. According to this representation of moral status, an AI system will have some moral standing if it is capable of experiencing qualia, including the ability to experience pain. A sentient AI system is more like a live animal than a stuffed toy or a wind-up doll, even if it lacks language and other higher cognitive abilities. Unless there are enough compelling moral justifications, it is unethical to cause suffering to a mouse. Any sentient AI system would have the same limitations. An AI system would have complete moral status, equal to that of humans, if it possessed consciousness as well as sapience of a type comparable to that of a normal adult human. Even if it lacks language and other higher cognitive abilities, a sentient AI system is more like a real animal than a stuffed toy or a wind-up doll. Inflicting pain on a mouse is immoral unless there are compelling moral justifications for doing so. Any sentient AI system would follow the same rules. An artificial intelligence system would have full moral status and be on par with humans if it has sapience that is comparable to that of a normal adult human. The Substrate NonDiscrimination Principle does not suggest that a digital computer could be aware or that it could act similarly to a person. Of course, a substance's moral significance might depend on how it affects functionality or sentience. But if we hold these things constant, it doesn't matter morally if a being is comprised of silicon or carbon, or whether its brain employs neurotransmitters or semi-conductors. Another premise that might be put out is that the moral standing of AI systems is not fundamentally affected by the fact that they are artificial, i.e., the result of intentional design. This might be stated as follows:

The principle of ontogeny non-discrimination states that two creatures have the same moral standing if they share the same functioning and awareness experience and only differ in how they came into being. Though in certain areas, especially in the past, the notion that one's moral standing depends on one's lineage or caste has been influential, this theory is now largely recognized in the human instance. We don't think that the moral standing of the offspring is necessarily affected

by causal factors like family planning, assisted birth, in vitro fertilization, gamete selection, intentional enhancement of maternal nutrition, etc., which introduce a deliberate choice and design element in the creation of human persons. Most people agree that if a human clone were to be born, it would have the same moral position as any other human newborn, including those who reject human reproductive cloning for ethical or religious grounds. This justification applies to the situation involving wholly constructed cognitive systems under the Principle of Ontogeny NonDiscrimination. Of course, it is conceivable for the conditions of creation to have an impact on the offspring that changes its moral standing. For instance, if a technique was carried out during conception or gestation that resulted in the development of a human child without a brain, this knowledge about ontogeny would be important to our evaluation of the moral standing of the offspring. However, the anencephalic infant would have the same moral standing as other anencephalic children of a similar kind, even those that were the result of completely natural processes. The qualitative difference between the two—the fact that one has a mind while the other does not—underlies the moral status difference between an anencephalic kid and a typical youngster. The Principle of Ontogeny Non-Discrimination does not apply since the two offspring do not have the same functioning and conscious experience. Although the Principle of Ontogeny Non-Discrimination claims that a being's ontogeny has no fundamental impact on its moral standing, it does not exclude the possibility that ontogenetic information may influence the obligations that different moral agents have toward the concerned being. Parents owe their children specific obligations that they do not owe to other children and that they would not owe even if another kid existed who was qualitatively similar to their own. The claim that the owners or creators of an AI system with moral status may have unique obligations to their artificial mind that they do not have to another artificial mind is also consistent with the Principle of Ontogeny Non-Discrimination, even if the minds in question are qualitatively similar and share the same moral standing. Many problems regarding how we should treat artificial minds may be resolved by applying the same moral standards that we use to evaluate our obligations in more familiar circumstances, providing the concepts of non-discrimination with regard to substrate and ontogeny are recognized. We should treat an artificial mind the same way we would treat a natural human mind in a similar circumstance, inasmuch as moral responsibilities derive from moral status concerns. The issue of creating an ethics for the treatment of artificial minds is made easier by this. Even if we adopt this position, we will still need to address a number of fresh ethical issues that the aforementioned principles do not address. Because artificial minds have the potential to differ greatly from those of regular humans or animals, new ethical issues might develop. We must take into account how these innovative qualities will impact artificial brains' moral standing and what it would entail to respect the moral standing of such unusual minds.

Conclusion

The move of AI algorithms toward more humanlike reasoning portends expected challenges, despite the fact that contemporary AI gives us few ethical problems that are not already present in the design of autos or power plants. AI algorithms may perform social functions, introducing new design criteria like predictability and openness. As a result, new types of safety assurance and the engineering of artificial ethical concerns may be necessary when sufficiently general AI algorithms no longer operate in predictable situations. AIs having sufficiently developed mental states, or the correct kind of states, will have moral standing, and some may even be considered to be persons—although they may be substantially different from those who already exist and maybe subject to other laws. Finally, the possibility of AIs possessing superhuman intelligence and strength forces us to face the very difficult task of developing an algorithm that produces superethical conduct. These problems may seem far-reaching, but it seems certain that we will face them, and there are plenty of ideas for current study areas in them.

References:

- Arkhipov, V., Naumov, V. On Certain Issues Regarding the Theoretical Grounds for Developing Legislation on Robotics: Aspects of Will and Legal Personality // Zakon. 2017, No. 5, p. 167.
- Arkhipov, V., Naumov, V. Op. cit., p. 164.
- Asaro P. The Liability Problem for Autonomous Artificial Agents // AAAI Symposium on Ethical and Moral Considerations in Non-Human Agents, Stanford University, Stanford, CA. March 21–23, 2016, p. 191.
- Asaro P. The Liability Problem for Autonomous Artificial Agents, p. 193.
- Asaro P., “From Mechanisms of Adaptation to Intelligence Amplifiers: The Philosophy of W. Ross Ashby,” in Wheeler M., Husbands P., and Holland O.

(eds.) The Mechanical Mind in History, Cambridge, MA: MIT Press: pp. 149–184

- Asimov, Isaac. 1942. "Runaround." Astounding Science-Fiction, March, 94–103.
- Bostrom, Nick, and Milan M. Ćirković, eds. 2008. Global Catastrophic Risks, New York: Oxford University Press. Chalmers, David John. 1996.
- Bostrom, Nick. 2002. "Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards." Journal of Evolution and Technology 9. http://www.jetpress.org/volume9/risks.html.
- Brożek B., Jakubiec M. On the Legal Responsibility of Autonomous Machines // Artificial Intelligence Law. 2017, No. 25(3), pp. 293–304.
- D. Edmonds, Would You Kill the Fat Man? The Trolley Problem and What Your Answer Tells Us About Right and Wrong, Princeton University Press, 2013.
- Goertzel, Ben, and Cassio Pennachin, eds. 2007. Artificial General Intelligence. Cognitive Technologies. Berlin: Springer. doi:10.1007/978-3-540-68677-4.
- Good, Irving John. 1965. "Speculations Concerning the First Ultraintelligent Machine." In Advances in Computers, edited by Franz L. Alt and Morris Rubinoff, 31–88. Vol. 6. New York: Academic Press. doi:10.1016/S0065-2458(08)60418-0.
- Hage J. Theoretical Foundations for the Responsibility of Autonomous Agents // Artificial Intelligence Law. 2017, No. 25(3), pp. 255–271.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2001. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. 1st ed. Springer Series in Statistics. New York: Springer.
- Hirschfeld, Lawrence A., and Susan A. Gelman, eds. 1994. Mapping the Mind: Domain Specificity in Cognition and Culture. New York: Cambridge University Press.
- Hofstadter, Douglas R. 2006. "Trying to Muse Rationally about the Singularity Scenario." Talk given at the Singularity Summit 2006, Stanford,
- Howard, Philip K. 1994. The Death of Common Sense: How Law is Suffocating America.New York: Random House.
- In countries that use the Anglo-Saxon legal system, the European Union and some Middle Eastern countries. This kind of liability also exists in certain former Soviet countries: Georgia, Kazakhstan, Moldova and Ukraine. It does not exist in Russia, although it is under discussion.
- Kamm, Frances M. 2007. Intricate Ethics: Rights, Responsibilities, and Permissible Harm. Oxford Ethics Series. New York: Oxford University Press. doi:10.1093/acprof:oso/9780195189698.001. 0001.
- Khanna V.S. Corporate Criminal Liability: What Purpose Does It Serve? // Harvard Law Review. 1996, No. 109, pp. 1477–1534.
- Kurzweil, Ray. 2005. The Singularity Is Near: When Humans Transcend Biology. New York: Viking.
- McDermott, Drew. 1976. "Artificial Intelligence Meets Natural Stupidity." SIGART Newsletter (57): 4–9. doi:10.1145/1045339.1045340.
- Sandberg, Anders. 1999. "The Physics of Information Processing Superobjects: Daily Life Among the Jupiter Brains." Journal of Evolution and Technology http://www.jetpress.org/volume5/ Brains2.pdf.
- U. Pagallo, The Laws of Robots. Crimes, Contracts, and Torts. Springer, 2013, p. 36.
- Vinge, Vernor. 1993. "The Coming Technological Singularity: How to Survive in the Post-Human Era." In Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace, 11–22. NASA Conference Publication 10129. NASA Lewis Research Center. http://ntrs.nasa.gov/archive/ nasa/casi.ntrs.nasa.gov/19940022855_1994022855.pdf.
- Wallach, Wendell, and Colin Allen. 2009. Moral Machines: Teaching Robots Right from Wrong.New York: Oxford University Press. doi:10.1093/acprof:oso/9780195374049.001.0001.
- Warren, Mary Anne. 1997. Moral Status: Obligations to Persons and Other Living Things.Issues in Biomedical Ethics. New York: Oxford University Press. doi:10.1093/acprof:oso/9780198250401.001.0001.
- Winkler A. We the Corporations: How American Businesses Won Their Civil Rights. Liverlight, 2018. See a description here: https://www.nytimes.com/2018/03/05/books/review/adam-winkler-we-the-corporations.html.