

Marcin GARBOWSKI
John Paul II Catholic University of Lublin
Methodology of Science Department at the Faculty of Philosophy
marcin.s.garbowski@gmail.com

A CRITICAL ANALYSIS OF THE ASILOMAR AI PRINCIPLES

Streszczenie. The article focuses on the analysis of a set of principles concerning the development of artificial intelligence, which were agreed upon at the Asilomar international conference in January 2017. The declaration was signed by leading scholar in the field of transhumanism and representatives of the hi-tech enterprises. The purpose of the article is to lay out the core assumptions behind this established set of norms that aim at the creation of so call „friendly AI” as well as their practical implications.

Słowa kluczowe: Artificial intelligence, philosophy of technology, transhumanism

KRYTYCZNA ANALIZA ZASAD DOTYCZĄCYCH SZTUCZNEJ INTELIGENCJI Z ASILOMAR

Abstract. Przedmiotem artykułu jest analiza zestawu zasad uzgodnionych na międzynarodowej konferencji w Asilomar ze stycznia 2017 roku dotycząca rozwoju sztucznej inteligencji. Sygnatariuszami porozumienia są czołowi badacze tematyki związanej z transhumanizmem, ale także przedstawiciele koncernów z sektora hi-tech. Celem artykułu jest przedstawienie założeń stojących za deklarowanym szkicem norm, mających na celu wytworzenie tzw. Przyjaznej sztucznej inteligencji oraz ich praktycznych konsekwencji.

Keywords: Sztuczna inteligencja, filozofia techniki, transhumanism

1. Introduction

When we consider the various technological breakthroughs which in the past decades have had the greatest impact on a broad scope of human activity and which are bound to increase it in the coming years, artificial intelligence certainly occupies a prominent place.

Defining this particular type of technology is not easy, yet one can attempt to include in its denotation a vast array of artifacts, the prime function of which is to aid or replace tasks performed by humans associated with “intelligent behavior” in a broad sense. Although with the use of technology humanity has been altering various aspects of performance by the use and inventions of ever more complex tools, it is customary that we speak of the dawn of “artificial intelligence” since the invention of computers in the twentieth century. Rapid advancements in telecommunication, analytics and other applications of simulated intelligent processes have led to the creation of the internet, personal computers, advanced diagnostics, automation in industry which continue to have an immense impact on our civilization worldwide as well as the daily lives of billions of people – solving many crucial problems and creating new ones. This spectacular progress has provoked philosophers, scientists and technological entrepreneurs to formulate hypotheses on where and how quickly extrapolations of these trends might lead us – ranging from promising visions of reaching a near God-like status to, conversely, catastrophic extinction scenarios.

One might say that applications of computing to technological artifacts do not require particular attention from the technological assessment community or that scenarios considered by the so called “transhumanists” are akin to science-fiction. And yet the subject of artificial intelligence has been receiving an increasing amount of attention in the media in recent years with such prolific figures like physicist Stephen Hawking¹ or technological entrepreneur and billionaire Elon Musk² emphasizing the risks connected with the rapid advancement in the field of artificial intelligence. Artificial intelligence by way of a variety of recent breakthroughs and increasing amount of investment from both governmental and corporate sources must be a central subject within the realm of technology assessment.

Although it is hard to grasp the abundance of the work performed by technological companies in the field of artificial intelligence, it is possible to analyze the guidelines which representatives of the industry as well as researchers in the field have accepted to follow as a part of a broad consensus. Such guidelines, called the “Asilomar AI Principles,” have been coined at a conference held at the Asilomar Conference Ground in Pacific Grove, California in January of 2017 – at the same place where an analogous set of values was accepted in February 1975 in reference to biotechnology³.

This paper shall present a brief introduction to the field of ethics of artificial intelligence and attempts to regulate the technology throughout a variety of fields, concentrating on the principles presented in the Asilomar declaration.

¹ Cellan-Jones R.: Stephen Hawking warns artificial intelligence could end mankind, <http://www.bbc.com/news/technology-30290540>.

² Gibbs S.: Elon Musk: artificial intelligence is our biggest existential threat, <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>.

³ Berg P.: Meetings that changed the world: Asilomar 1975: DNA modification secured. “Nature”, Vol. 455, p. 290-291, <https://www.nature.com/nature/journal/v455/n7211/full/455290a.html>.

2. AI ethics: sources and significance

The main sources of inspiration for the statements constituting the body of artificial intelligence ethics are the following:

- Science fiction – both literature and film, and more recently computer and video games – has been providing inspiration for researchers in the field of AI Ethics. Fictional accounts of the application of high level AI for both malevolent and benevolent uses have been present in the genre since the beginning, as far back as the 1950s. The science-fiction writer Isaac Asimov is the author of the often-cited “laws of robotics,” which he first included in his collection of short stories *I, robot*:

- A robot may not injure a human being or, through inaction, allow a human being to come to harm.
- A robot must obey orders given it by human beings except where such orders would conflict with the First Law.
- A robot must protect its own existence as long as such protection does not conflict with the First or Second Law⁴.

Relations between humans and AI are the subject matter of novels and movies such as Phillip K. Dick’s *Do Androids Dream of Electric Sheep?*⁵, Stanisław Lem’s *Summa Technologiae*⁶ and *Cyberiad*⁷, James Cameron’s *The Terminator*⁸ series or the Wachowski sisters’ *The Matrix*⁹. Themes present in these and many other works can be treated as a point of departure for reflections in the field of philosophy.

- Philosophy and futurology – the topic of AI, its development and effect on the future of humanity – has been present in the works of visionaries ever since the dawn of the computer era. One of the early specialists in the field, Irving J. Good, wrote of a coming “intelligence explosion” in his 1965 essay¹⁰. Marvin Minsky set forth the positive sides of the coming progress of AI in *The Society of Mind*¹¹, as did Ray Kurzweil in *The Age of Intelligent Machines* and subsequent books, where he described the concept of the “technological singularity”¹². This concept received a large deal of attention among a wider audience with the rapid process of computerization and the proliferation of the internet since the 1990s. This hypothetical

⁴ Asimov I.: *I, robot*. Gnome Press, New York 1950, p. 26, https://www.ttu.ee/public/m/mart-murdvee/Techno-Psy/Isaac_Asimov_-_I_Robot.pdf.

⁵ Dick P.K.: *Do Androids Dream of Electric Sheep?* Doubleday Press, New York 1968.

⁶ Lem S.: *Summa Technologiae*. University of Minnesota Press, Minneapolis-London 2013.

⁷ Lem S.: *The Cyberiad*. Harcourt, Brace & Company, New York 1985.

⁸ Cameron J.: *The Terminator*. USA 1984.

⁹ Wachowski L., Wachowski L.: *The Matrix*. USA 1999.

¹⁰ Good I.J.: *Speculations Concerning the First Ultra-intelligent Machine*. “*Advances in Computers*”, Vol. 6, 1965.

¹¹ Minsky M.: *The Society of Mind*. Simon and Schuster, New York 1987.

¹² Kurzweil R.: *The Age of Intelligent Machines*. MIT Press, Boston 1990.

event, when humanity merges with a highly developed form of AI, has generated an entire intellectual movement whose members describe themselves as transhumanists, but also provoked criticism from various circles. Even among people connected with the movement or the development of these technologies per se, concerns as well as visions of potential threats are abundant. What is more, well-known enthusiasts of technological progress, such as Nick Bostrom, mention the development of malevolent super-intelligent AI as an existential risk¹³; James Hughes describes the consequences of cyborgization on politics and society¹⁴ and Hugo de Garis describes a potential conflict, which may lead to the obliteration of the human race in a Terminatoresque scenario¹⁵.

- Corporate and governmental policy – in recent years the topic of regulating human vs. AI relations has reached beyond academic milieus and the perspective of creating a super-intelligent threat is perceived as a challenge by, among others, Elon Musk, who considers the task of developing benevolent AI crucial to preventing catastrophic scenarios¹⁶. Projects connected with the development of such technologies have been undertaken by major private and public sector entities in the US¹⁷, Russia¹⁸, China¹⁹ and the EU²⁰.
- Legislature – initial actions in the field of legally regulating the status of AI have been undertaken in the European Parliament at the beginning of 2017²¹. Yet currently the main concern is not an existential risk for humanity in general, but the possibility of increasing unemployment caused by the fact that a number of professions may become obsolete as a result of the progress in the automation of many repetitive tasks. The concept of these proposed regulations is focused on the taxation of the output robots as “legal persons” and the distribution of revenue among the overall population in the form of “Universal Basic Income”²². The proposed legislation also calls for the

¹³ Kurzweil R.: *The Age of Spiritual Machines*. Viking Press, New York 1999.

¹⁴ Kurzweil R.: *The Singularity is Near*. Viking Press, New York 2006.

¹⁵ Bostrom N.: *Future of Humanity*, [in:] Olsen J.K.B., Selinger E., Riis S.: *New Waves in Philosophy of Technology*. Palgrave MacMillan, New York 2009.

¹⁶ Hughes J.: *Citizen Cyborg: Why Democratic Societies Must Respond to the Redesigned Human of the Future*. Westview Press, 2004.

¹⁷ de Garis H.: *The Artilect War: Cosmists vs. Terrans: A Bitter Controversy Concerning Whether Humanity Should Build Godlike Massively Intelligent Machines*. ETC Publications, 2005.

¹⁸ Mack E.: *Why Elon Musk Spent \$10 Million To Keep Artificial Intelligence Friendly*, <https://www.forbes.com/sites/ericmack/2015/01/15/elon-musk-puts-down-10-million-to-fight-skynet/#4350a74b2e5b>.

¹⁹ DARPA (Defense Advanced Research Projects Agency), <http://www.darpa.mil/about-us/darpa-perspective-on-ai> and the Google corporation (with Ray Kurzweil as the supervisor of the engineering department, https://research.google.com/pubs/Machine_Intelligence.html are engaged in projects connected with the development of AI).

²⁰ Project Avatar founded by Russian billionaire Dmitri Yetskov, <http://2045.com/>.

²¹ The China Brain Project, <http://www.irma-international.org/viewtitle/46407/>.

²² The Human Brain Project, <http://www.humanbrainproject.eu/en/>.

installation of a “kill-switch” or another device that may de-activate a robot, when it potentially gets out of hand²³.

The reasons why we consider artificial intelligence as a significant topic for ethical analysis are manifold and may be summarized in the following categories:

- Increasing social and economic impact: the increasing level of automation in various segments of activities associated with human labor as well as in military applications is provoking long-term ominous predictions, which foresee threats to the global social order or even to the existence of humanity itself.
- Strategic significance: in any far reaching strategic planning connected with economic and social foresight one must include a SWOT analysis referring to the positive and negative aspects of proliferation of artificial intelligence in various fields of human activity.
- Immense allocation of resources by the public and private sector: projects connected with the development of various branches of industry and services applying inventions using AI are funded by the largest corporations in such fields as the military (Unmanned Aviation Vehicles, so called “drones”)²⁴, big data²⁵ production²⁶, transportation²⁷, entertainment²⁸ and even agriculture²⁹.
- Broad academic interest: the problem of the social and moral consequences of the development and proliferation of AI has become the topic of numerous dissertations and conferences in recent years, and special institutes and even universities have been established in this realm of interest, e.g. The Future of Humanity Institute at Oxford University³⁰, the Singularity University at Silicon Valley³¹ or the Centre for the Study of Existential Risk at Cambridge University³².

²³ Robots: Legal Affairs Committee calls for EU-wide rules, <http://www.europarl.europa.eu/news/en/news-room/20170110IPR57613/robots-legal-affairs-committee-calls-for-eu-wide-rules>.

²⁴ Bulman M.: EU Universal Income must be 'seriously considered' as rise of robots threatens mass unemployment, <http://www.independent.co.uk/life-style/gadgets-and-tech/universal-basic-income-robots-eu-meps-unemployment-mady-delvaux-stehres-european-parliament-a7527661.html>.

Future of Humanity Institute at Oxford University, <https://www.fhi.ox.ac.uk/>.

²⁵ Kottasova I.: Europe calls for mandatory 'kill switches' on robots, <http://money.cnn.com/2017/01/12/technology/robot-law-killer-switch-taxes/index.html>.

²⁶ Kopstein J.: A.I. To Make US War Drones Deadlier, Less Accountable, <http://www.vocativ.com/430867/artificial-intelligence-algorithmic-warfare-drone-strikes/>.

²⁷ Canton J.: From Big Data to Artificial Intelligence: The Next Digital Disruption, http://www.huffingtonpost.com/james-canton/from-big-data-to-artifici_b_10817892.html.

²⁸ Wilson K.: Age of the robots, http://usa.chinadaily.com.cn/business/2017-05/19/content_29423831.htm.

²⁹ Pesce M.: Autonomous cars are about to do to transport what the internet did to information, https://www.theregister.co.uk/2017/02/27/autonomous_cars_2040/.

³⁰ Kriegsman M.: Augmented reality: An enterprise business imperative, <http://www.zdnet.com/article/augmented-reality-an-enterprise-business-imperative/>.

³¹ Daniels J.: Future of farming: Driverless tractors, ag robots, <http://www.cnn.com/2016/09/16/future-of-farming-driverless-tractors-ag-robots.html>.

³² Future of Humanity Institute at Oxford University, <https://www.fhi.ox.ac.uk/>.

3. The Asilomar AI principles

The effect of the above-mentioned concerns was the above mentioned Asilomar conference, where the key figures dealing with AI issues from a business and academic perspective joined forces to create a joint declaration³³. So far it has been signed by 1200 researchers in the field of AI as well as more than 2300 other endorsers, among others Elon Musk, Ray Kurzweil, Stephen Hawking and David Chalmers, to mention the more notable thinkers and entrepreneurs.

The declaration consists of 23 points divided into 3 sections entitled “Research Issues,” “Ethics and Values” and “Longer-term Issues”. I present the entire text of these principles below as terminology linked with the axiological aspects of AI is intertwined with technical hints:

Research Issues

- 1) Research Goal: The goal of AI research should be to create not undirected intelligence, but beneficial intelligence.
- 2) Research Funding: Investments in AI should be accompanied by funding for research on ensuring its beneficial use, including thorny questions in computer science, economics, law, ethics, and social studies, such as:
 - How can we make future AI systems highly robust, so that they do what we want without malfunctioning or getting hacked?
 - How can we increase our prosperity through automation while maintaining people’s resources and purpose?
 - How can we update our legal systems to be more fair and efficient, to keep pace with AI, and to manage the risks associated with AI?
 - What set of values should AI be aligned with, and what legal and ethical status should it have?
- 3) Science-Policy Link: There should be a constructive and healthy exchange between AI researchers and policy-makers.
- 4) Research Culture: A culture of cooperation, trust, and transparency should be fostered among researchers and developers of AI.
- 5) Race Avoidance: Teams developing AI systems should actively cooperate to avoid corner-cutting on safety standards”.

Ethics and Values

- 6) Safety: AI systems should be safe and secure throughout their operational lifetime, and verifiably so where applicable and feasible.

³³ Singularity University, <https://su.org/>.

- 7) Failure Transparency: If an AI system causes harm, it should be possible to ascertain why.
- 8) Judicial Transparency: Any involvement by an autonomous system in judicial decision-making should provide a satisfactory explanation auditable by a competent human authority.
- 9) Responsibility: Designers and builders of advanced AI systems are stakeholders in the moral implications of their use, misuse, and actions, with a responsibility and opportunity to shape those implications.
- 10) Value Alignment: Highly autonomous AI systems should be designed so that their goals and behaviors can be assured to align with human values throughout their operation.
- 11) Human Values: AI systems should be designed and operated so as to be compatible with ideals of human dignity, rights, freedoms, and cultural diversity.
- 12) Personal Privacy: People should have the right to access, manage and control the data they generate, given AI systems' power to analyze and utilize that data.
- 13) Liberty and Privacy: The application of AI to personal data must not unreasonably curtail people's real or perceived liberty.
- 14) Shared Benefit: AI technologies should benefit and empower as many people as possible.
- 15) Shared Prosperity: The economic prosperity created by AI should be shared broadly, to benefit all of humanity.
- 16) Human Control: Humans should choose how and whether to delegate decisions to AI systems, to accomplish human-chosen objectives.
- 17) Non-subversion: The power conferred by control of highly advanced AI systems should respect and improve, rather than subvert, the social and civic processes on which the health of society depends.
- 18) AI Arms Race: An arms race in lethal autonomous weapons should be avoided.

Longer-term Issues

- 19) Capability Caution: There being no consensus, we should avoid strong assumptions regarding upper limits on future AI capabilities.
- 20) Importance: Advanced AI could represent a profound change in the history of life on Earth, and should be planned for and managed with commensurate care and resources.
- 21) Risks: Risks posed by AI systems, especially catastrophic or existential risks, must be subject to planning and mitigation efforts commensurate with their expected impact.
- 22) Recursive Self-Improvement: AI systems designed to recursively self-improve or self-replicate in a manner that could lead to rapidly increasing quality or quantity must be subject to strict safety and control measures.

23) Common Good: Superintelligence should only be developed in the service of widely shared ethical ideals, and for the benefit of all humanity rather than one state or organization”³⁴.

As one can notice many of the terms included in this list are not clearly defined and, moreover, can be understood in a variety of ways. The very concept of what is “beneficial” or what one can consider as a “common good” is linked with a notion which has been debated throughout history of philosophy and various schools of thought and is still discussed today. Apart from the philosophical views expressed elsewhere by some of the more notable signatories of the document there is no hint at a particular axiological system.

Although this is not stated explicitly within the contents of these principles, the direction, in which the scientific community is heading in its reflection on the future of AI is akin to that of classical philosophy, where the fulfillment of a particular good is connected with a particular purpose (Greek *telos*). Although it would seem that advocates of a naturalistic or even a Darwinian approach to technological development should trust that the spontaneous, “undirected” process of creating intelligence should deliver the most desirable effects, the vision of a self-aware AI that could consider humanity redundant brings back a more Aristotelian perspective on technology. The AI in order to be beneficial, and therefore not autonomous from the human being’s control – but a tool, must fulfil its purpose. This purpose cannot be located in an axiological void – it has to be consistent with an entire set of values which are ultimately metaphysical.

In order to determine what is “good” for humanity in relation to AI, one cannot omit addressing, at least, these fundamental questions shall make this discussion clearer:

- What is a “human being”?
- What is “intelligence”?
- What is the difference between “human” and “machine intelligence”?
- Is human flourishing dependent on the amount of intelligence?
- How can we estimate the desired level of “common good”?

In order to address these questions one might to dissect the guidelines suggested in the declaration according to a pattern which I present below:

- Methodology: Defining core concepts crucial for the AI sector; meanings of the fundamental terms must not be implied and any element of doubt must be eviscerated.
- Ethics: Creating consistent value systems; values based on these axioms must be articulated in a way that can be easily transposed to the language used to program AI machines.
- Law: Adapting ethical codes and legislature to the established value systems; the execution of such codes must be backed by effective means, which can supervise a practical application of the guidelines in the AI industry.

³⁴ Center for the Study of Existential Risk, <http://cser.org/>.

- Application and assessment: Independent institutions must oversee the way these regulations are affecting the practical output of these technologies and provide feedback to the industry.

It might turn out that creating a unified paradigm for assessing artificial intelligence or in fact any technology affecting the human condition in a more profound way is impossible. Various questions may receive different answers depending on what metaphysical or anthropological base we assume. Defining the core notions that form the framework for our assessment may differ immensely when we take on a naturalistic or else a theistic viewpoint. Various axiological viewpoints may be applied to define what is “right” for a machine to do. When speaking of “beneficial AI” the fundamental value we are discussing is that of the “good” and in this specific case we mean the good of a human being or humanity as a whole.

The scientific method has no definite answers in itself to transcendental values associated with what is “good” and even those who advocate a naturalistic worldview must assume an axiomatic set of suppositions when it comes to morality. Undoubtedly it is possible to revert to moral relativism, but would we then consider a morally relativistic super-intelligent AI beneficial towards human beings, if it perceived the concept of human “good” in the way human moral relativists treated it in the past?

4. Conclusion

Undoubtedly the problems mentioned in Asilomar AI Principles declaration are actual challenges that must be addressed. The principles presented in the declaration are a good point of departure for further discussion and analysis, but without making the basic ontological criteria more precise and backed by strict regulations these points are not enough to secure a safe relationship between humans and AI. The values, which humanity may have to assume in its endeavor to prevent the proliferation of malevolent AI may turn out to be much closer to an objective truth rather than a democratic consensus of “widely shared ethical ideals,” which in itself may be shocking for proponents of materialistic philosophy. Whether this truth about the value of humanity and flourishing is accepted as the Kierkegaardian leap of faith or another approach is chosen, it may happen that the human species encountering a possible intellectual peer and threat to its existence shall embrace an objective and transcendental concept of the good life.

Bibliography

1. Cellan-Jones R.: Stephen Hawking warns artificial intelligence could end mankind, <http://www.bbc.com/news/technology-30290540>.
2. Gibbs S.: Elon Musk: artificial intelligence is our biggest existential threat, <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>.
3. Berg P.: Meetings that changed the world: Asilomar 1975: DNA modification secured. "Nature", Vol. 455, p. 290-291, <https://www.nature.com/nature/journal/v455/n7211/full/455290a.html>.
4. Asimov I.: I, robot. Gnome Press, New York 1950, p. 26, https://www.ttu.ee/public/m/mart-murdvee/Techno-Psy/Isaac_Asimov_-_I_Robot.pdf.
5. Dick P.K.: Do Androids Dream of Electric Sheep? Doubleday Press, New York 1968.
6. Lem S.: Summa Technologiae. University of Minnesota Press, Minneapolis-London 2013.
7. Lem S.: The Cyberiad. Harcourt, Brace & Company, New York 1985.
8. Cameron J.: The Terminator. USA 1984.
9. Wachowski L., Wachowski L.: The Matrix. USA 1999.
10. Good I.J.: Speculations Concerning the First Ultra-intelligent Machine. "Advances in Computers", Vol. 6, 1965.
11. Minsky M.: The Society of Mind. Simon and Schuster, New York 1987.
12. Kurzweil R.: The Age of Intelligent Machines. MIT Press, Boston 1990
13. Kurzweil R.: The Age of Spiritual Machines. Viking Press, New York 1999.
14. Kurzweil R.: The Singularity is Near. Viking Press, New York 2006.
15. Bostrom N.: Future of Humanity, [in:] Olsen J.K.B., Selinger E., Riis S.: New Waves in Philosophy of Technology. Palgrave MacMillan, New York 2009.
16. Hughes J.: Citizen Cyborg: Why Democratic Societies Must Respond to the Redesigned Human of the Future. Westview Press, 2004.
17. de Garis H.: The Artilect War: Cosmists vs. Terrans: A Bitter Controversy Concerning Whether Humanity Should Build Godlike Massively Intelligent Machines. ETC Publications, 2005.
18. Mack E.: Why Elon Musk Spent \$10 Million To Keep Artificial Intelligence Friendly, <https://www.forbes.com/sites/ericmack/2015/01/15/elon-musk-puts-down-10-million-to-fight-skynet/#4350a74b2e5b>.
19. DARPA (Defense Advanced Research Projects Agency), <http://www.darpa.mil/about-us/darpa-perspective-on-ai> and the Google corporation (with Ray Kurzweil as the supervisor of the engineering department, <https://research.google.com/pubs/MachineIntelligence.html> are engaged in projects connected with the development of AI).
20. Project Avatar founded by Russian billionaire Dmitri Yetskov, <http://2045.com/>.

21. The China Brain Project, <http://www.irma-international.org/viewtitle/46407/>.
22. The Human Brain Project, <http://www.humanbrainproject.eu/en/>.
23. Robots: Legal Affairs Committee calls for EU-wide rules (press release), <http://www.europarl.europa.eu/news/en/news-room/20170110IPR57613/robots-legal-affairs-committee-calls-for-eu-wide-rules>.
24. Bulman M.: EU Universal Income must be 'seriously considered' as rise of robots threatens mass unemployment, <http://www.independent.co.uk/life-style/gadgets-and-tech/universal-basic-income-robots-eu-meps-unemployment-mady-delvaux-stehres-european-parliament-a7527661.html>.
25. Kottasova I.: Europe calls for mandatory 'kill switches' on robots, <http://money.cnn.com/2017/01/12/technology/robot-law-killer-switch-taxes/index.html>.
26. Kopstein J.: A.I. To Make US War Drones Deadlier, Less Accountable, <http://www.vocativ.com/430867/artificial-intelligence-algorithmic-warfare-drone-strikes/>.
27. Canton J.: From Big Data to Artificial Intelligence: The Next Digital Disruption, http://www.huffingtonpost.com/james-canton/from-big-data-to-artifici_b_10817892.html.
28. Wilson K.: Age of the robots, http://usa.chinadaily.com.cn/business/2017-05/19/content_29423831.htm.
29. Pesce M.: Autonomous cars are about to do to transport what the internet did to information, https://www.theregister.co.uk/2017/02/27/autonomous_cars_2040/.
30. Kriegsman M.: Augmented reality: An enterprise business imperative, <http://www.zdnet.com/article/augmented-reality-an-enterprise-business-imperative/>.
31. Daniels J.: Future of farming: Driverless tractors, ag robots, <http://www.cnbc.com/2016/09/16/future-of-farming-driverless-tractors-ag-robots.html>.
32. Future of Humanity Institute at Oxford University, <https://www.fhi.ox.ac.uk/>.
33. Singularity University, <https://su.org/>.
34. Center for the Study of Existential Risk, <http://cser.org/>.
35. A Principled AI Discussion in Asilomar. Future of Life Institute, <https://futureoflife.org/2017/01/17/principled-ai-discussion-asilomar/>.
36. Asilomar AI Principles, <https://futureoflife.org/ai-principles/>.