

## PART I

# **AI – Humans vs. Machines**



## CHAPTER 2

# Artificial Intelligence (AI): When Humans and Machines Might Have to Coexist

Andreas Kaplan

### Introduction

Artificial intelligence (AI), defined as a ‘system’s ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation’ (Kaplan and Haenlein 2019, 17), will likely have a deep impact on human beings and society at large. The recent COVID-19 pandemic has particularly accelerated and accentuated society’s digitalisation and strongly influences the future relationship between human beings and AI-driven machines (Haenlein and Kaplan 2021).

Various opinions and viewpoints on the future altered by advances in AI exist, ranging from horror scenarios as stated by Tesla CEO Elon Musk, to utopian scenarios like the vision of Google Chief Engineer Raymond Kurzweil. While Musk fears that AI might lead to nothing less than a third world war, Kurzweil believes that AI will enhance humans instead of replacing them. Expressing these opposing views, in 2018, theoretical physicist Stephen Hawking proclaimed that AI can ‘either be the best, or the worst thing, ever to happen to humanity’.

---

#### How to cite this book chapter:

Kaplan, A. 2021. Artificial Intelligence (AI): When Humans and Machines Might Have to Coexist. In: Verdegem, P. (ed.) *AI for Everyone? Critical Perspectives*. Pp. 21–32. London: University of Westminster Press. DOI: <https://doi.org/10.16997/book55.b>. License: CC-BY-NC-ND 4.0

Clearly, humans will need to coexist with machines. Jobs traditionally done by humans will be shifted towards AI systems. Artificial intelligence is already able to translate languages, diagnose illnesses, assist in retail (Kaplan 2020c), and the like – in several cases, better than the human workforce. Human jobs might be created in the future that are unimaginable now, similar to the fact that nobody really predicted the job of mobile app designers just a few years ago.

In this world, AI would rather be augmenting and complementing – rather than replacing – humans in their work. In the pessimistic case, i.e., massive unemployment, ideas such as universal basic income are already being discussed. Fundamental philosophical questions would need to be answered surrounding life for humans when most of our work is done by AI systems. In any case, the State will certainly have to come up with a set of rules governing this human < > machine coexistence and interdependence. Society overall is thus challenged.

This chapter has a look at artificial intelligence, its history and its evolutionary stages. Furthermore, what challenges might arise in the future when humans will have to learn to live among machines and robots will be discussed. This will be done by analysing challenges concerning algorithms and organisations, challenges with respect to (un)employment, and looking at democracy and freedom potentially jeopardised due to AI progress.

## Artificial Intelligence: Definition and Classification

Artificial intelligence is a rather fuzzy concept, and quite difficult to define. At least two reasons can be proposed for the difficulty in formulating a definition therefore: firstly, it is not easy to find a clear definition for what intelligence in general is, as it depends largely upon the context. Thus intelligence is described in several different ways such as the capacity for learning, reasoning, planning, understanding, critical thinking, creativity, and last but not least, problem solving.

Secondly, artificial intelligence is a moving target: advances previously considered to AI with time will not be considered as such as soon as we get used to them. This phenomenon is known as the *AI effect*. As McCordick (2004, 204) formulated it: ‘It’s part of the history of the field of artificial intelligence that every time somebody figured out how to make a computer do something – play good checkers, solve simple but relatively informal problems – there was chorus of critics to say, “that’s not thinking.”’ Or as Rodney Brooks, MIT’s Artificial Intelligence Laboratory director, explains, ‘Every time we figure out a piece of it, it stops being magical; we say, “Oh, that’s just a computation”, and will not count as artificial intelligence any longer’ (Kahn 2002).

One of the prevailing definitions of artificial intelligence, as aforementioned, characterises AI as ‘a system’s ability to correctly interpret external data, to learn from such data and to use those learnings to achieve specific goals and tasks through flexible adaptation’ (Kaplan and Haenlein 2019, 17). Several further

definitions exist and experts disagree on how to best characterise artificial intelligence. By analysing different AI definitions, Russell and Norvig (2016), e.g., concluded that there are four main approaches for defining AI, i.e., see it as systems that (1) think like humans, (2) act like humans, (3) think rationally and (4) act rationally.

Often terms such as big data, machine learning or the Internet-of-Things (IoT) are incorrectly applied as synonyms for artificial intelligence, yet they are indeed differing concepts and terms. An AI-driven system needs big data from which to learn, which essentially are ‘datasets made up by huge quantities (volume) of frequently updated data (velocity) in various formats, such as numeric, textual or images/videos (variety)’ (Kaplan and Haenlein 2019, 17). Again, a variety of different definitions for big data exists: while one group of them focuses on what big data is, a second group stresses what big data actually does (Gandomi and Haider 2015). Such big data sets can derive from an organisation’s internal databases, third-party data or social media applications (Kaplan 2012; Kaplan and Haenlein 2010b).

Another possibility for obtaining big data is via the Internet-of-Things (Krotov 2017; Saarikko, Westergren and Blomquist 2017), which basically is an extension of internet connectivity into physical devices and everyday objects such as a refrigerator or a heater, equipped with sensors and software to collect and exchange data.

Machine learning, simply put, is ‘methods that help computers learn without being explicitly programmed’ (Kaplan and Haenlein 2019, 17), and is applied in order to identify underlying patterns within the big data, and as such is an essential element of artificial intelligence. A more elaborated definition comes from Mitchell (1997, 2) stating ‘A computer program is said to learn from experience  $E$  with respect to some class of tasks  $T$  and performance measure  $P$  if its performance at tasks in  $T$ , as measured by  $P$ , improves with experience  $E$ ’. AI is much broader than machine learning, as it additionally comprises such abilities as the perception of data (e.g., voice/image recognition, natural language processing, etc.) or the control and movement of objects (robotics or cybernetics).

Artificial intelligence can be classified into three types of systems: analytical, human-inspired and humanised (Kaplan and Haenlein 2019). *Analytical* AI contains characteristics consistent with cognitive intelligence only: generating cognitive representation of the world and using learning based on past experience to inform future decisions. *Human-inspired* AI contains elements of cognitive and emotional intelligence: understanding human emotions, in addition to cognitive elements, and considering them in their decision-making. *Humanised* AI contains characteristics of all types of competencies (i.e., cognitive, emotional and social intelligence), is able to be self-conscious, and is self-aware in interactions with others.

A robot driven by analytical artificial intelligence would be capable of answering queries concerning restaurant recommendations based on certain

objective characteristics. Human-inspired AI robots could additionally read a human's emotional state via facial recognition or tone of voice, and adapt its suggestions, e.g., a human who appears sad or depressed would not enjoy a restaurant with a lively atmosphere, whereas a happy human might totally enjoy such an environment. Finally, a humanised robot would understand when it was appropriate for it to offer to accompany the human or whenever this would not be appreciated, e.g., a couple insanely in love who would rather spend the time in intimate togetherness.

Finally, we must distinguish AI on the lower spectrum from so-called expert systems, often wrongly associated with artificial intelligence, as well as on the higher spectrum from skills that remain only possible for human beings: Expert systems are 'collections of rules programmed by humans in the form of if > then statements' (Kaplan and Haenlein 2019, p. 18). As these systems lack the ability to learn autonomously from external data, they should definitely not be counted as AI. Expert systems reconstruct human intelligence in a top-down manner (also called the knowledge-based or symbolic approach), considering that it can be codified as a set of predefined rules. In contrast, AI applies a bottom-up approach (also called the behaviour-based or connectionist approach) and imitates a brain's set-up (e.g., through neural networks) by using large quantities of data to infer knowledge independently.

The question that arises is what will remain human in the future and what cannot be imitated by AI systems, which is quite a tough question to answer. Most likely, humans will always have exclusivity when it comes to artistic creativity, Albert Einstein having pointed out that 'creativity is intelligence having fun'. Currently, it seems very improbable that AI systems will be able to be truly creative. But then again, the question is what exactly true creativity is, and who will be the judge of it?

### Artificial Intelligence: History and Evolution

To structure AI's history, we'll use an analogy of the four seasons: spring, summer, autumn and winter (Haenlein and Kaplan 2019). AI's birth period, i.e., spring, took place both in fiction as well as non-fiction. Regarding the former, Isaac Asimov, an American writer and professor of biochemistry at Boston University, published 'Runaround', a story revolving around an AI-driven robot, in 1942. In this story, Asimov's (1950, 40) three laws of robotics explicitly appear for the first time:

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

These three laws already hint at the difficulty of humans and robots coexisting. In any case, the robot in Asimov's story freezes in a loop of repetitive behaviour, as it doesn't find a solution for obeying laws 2 and 3 at the same time. 'Runaround' is therefore a cornerstone in the history of artificial intelligence, as it inspired generations of academics and researchers in the domain of AI.

Regarding the real world, we can refer to computer scientist Alan Turing's seminal paper 'Computing Machinery and Intelligence', published in 1950. Therein, Turing describes what now is known as the Turing test, or a test of a machine's ability to exhibit intelligent behaviour equivalent to, or indistinguishable from, that of a human. AI spring's climax can be pinpointed to the 1956, when Marvin Minsky and John McCarthy organised the Dartmouth Summer Research Project on Artificial Intelligence (DSRPAI) at Dartmouth College. It was at this workshop that the term *artificial intelligence* was coined.

After spring, there followed a couple of hot AI summers and very cold AI winters. While AI summers were characterised by huge enthusiasm and financing of AI, winters were marked by reduced funding and interest in artificial intelligence research. The first summer period lasted nearly 20 years. One of its successes was certainly ELIZA: Developed in 1966 by German-American computer scientist Joseph Weizenbaum, a professor at MIT, this computer program was so good at conversing with a human being that it appeared to pass the aforementioned Turing test.

General hype around AI and its development followed. However, this hype was soon replaced by disappointment and disenchantment. AI winter somehow had already begun when Marvin Minsky supposedly still contended that artificial intelligence could attain a human being's general average intelligence within three to eight years from that moment (Darrach 1970). As we all know, this did not occur; AI funding was heavily reduced and another AI summer did not happen until the 1980s, when the Japanese government decided to massively invest in AI and consequently the US DARPA followed. Success again was scarce, and summer was again followed by another cold winter.

We might have reached AI's autumn, completing the four seasons of artificial intelligence (Haenlein and Kaplan 2019), as a result of computational strength having constantly increased over recent years, rendering deep learning and artificial neural networks possible (Libai et al. 2020). This new era of AI is said to have begun in 2015 when AlphaGo, a computer program designed by Google, beat a (human) world champion in the Chinese board game Go. This event made the news around the world, and regenerated hype around the domain of artificial intelligence.

This hype might continue for quite some time, as we are currently only experiencing so-called first-generation AI applications, usually referred to as artificial narrow intelligence (ANI). Within such systems, AI is only applied to very specific tasks such as choosing which news items it will tell an individual during his or her morning before-work routine based on the individual's intellectual preferences.

Second-generation AI applications will be able to plan, solve and reason problems independently, even for actions for which they have not been programmed initially. Such artificial general intelligence (AGI) will thus be able to broaden its horizons autonomously, entering new areas and domains. For example, an AGI-powered system could, on top of conveying news headlines during one's morning routine, also learn to make coffee for the aforementioned individual preparing for work.

Finally, we might potentially even experience artificial super intelligence (ASI), the third generation of AI. Such truly self-conscious and self-aware AI systems, outperforming humans in (nearly) all domains, capable of general wisdom, scientific creativity and social skills, could render human beings redundant. As such, in our above example, the individual would not need to prepare for work anymore, as this could be done entirely by the ASI-powered machine or robot (Kaplan and Haenlein 2019). For a detailed discussion on the evolution of AI systems, we refer to Huang and Rust (2018).

## **Artificial Intelligence: Machines and Humans**

In the future, artificial intelligence will raise several challenges, and humans will have to learn to coexist with machines and robots. Pushed by the global COVID-19 health crisis, it is clear that AI will deeply impact societies around the world (Kaplan 2021). We will discuss some of these questions, looking at challenges in terms of algorithms and individual organisations; the employment market; and last but not least, democracy and human freedom potentially at stake due to advances in AI.

### *About Algorithms and Organisations*

When machines and humans coexist, it is important that both do what they are good at. As an illustration, let's have a look at a study by researchers from MIT's Computer Science and Artificial Intelligence Laboratory in cooperation with the machine-learning startup PatternEx (Conner-Simons 2016). AI systems and humans scored far better in identifying cyber-attacks when collaborating than when trying to do so separately. While the AI systems could crawl through enormous quantities of big data, humans were better at detecting anomalies, playing those back into the system. This iterative and collaborative approach was optimal.

Also, humans are better in behaving ethically and morally, while algorithms have problems doing so, as the notion of ethics and morals is difficult to program. Machines, however, are better at, e.g., utilitarian, repetitive tasks. While most humans would not consciously discriminate another individual for gender, sexual orientation, social background, or race, machines, not having a

conscience, are more likely to be biased, essentially because the data on which they were trained was biased. A study by Wilson, Hoffman and Morgenstern (2020) illustrates that several decision-support systems applied by judges may be racially biased (as a result of past rulings); and self-driving cars better detected lighter skin than darker tones, since their algorithm was trained using pictures among which were few people of colour.

Regulation and guidance is definitely needed in order to avoid such bias, to establish a good foundation for machine < > human collaboration. The development of specific requirements with respect to the testing and training of AI is likely the preferred approach, as opposed to regulating artificial intelligence itself. In addition, we could require AI warranties, consistent with safety testing in the case of physical goods. Thus, AI regulation could be stable over time even if the actual aspects of AI technology change (Kaplan and Haenlein 2020).

### *About (Un)Employment*

A tough challenge when human beings coexist with machines might be the evolution of the job market. Already, automation in manufacturing has led to a significant decrease in blue-collar jobs; advances in AI could lead to a similar decrease in white-collar jobs. AI systems already outperform medications in the identification of skin cancer and other tasks (Welch 2018).

For the moment, it appears that the time gain through AI's application is used for other tasks within the job, and does not necessarily lead to a human being's replacement. The Swedish bank, SEB, e.g., developed AIDA, an AI-driven virtual assistant responding to a vast range of customers' queries, such as how to make overseas payments or how to proceed when opening a bank account. AIDA is even capable of detecting a customer's mood by the tone of her or his voice and adapting its recommendations and suggestions thereto. In around 30% of situations, AIDA is not able to respond or help. In this case, the customer is transferred to a human. AIDA's implementation freed up human employees' time, which they then use for more complex demands, i.e., the 30% that exceeded AIDA's limitations.

A study by Wilson and Daugherty (2018, 117) suggested that it is in companies' interest not to replace employees with AI, as this would not be a long-term strategy. Looking at 1,500 corporations, they identified the best improvements in performance when machines and human beings work together, and concluded: 'Through such collaborative intelligence, humans and AI actively enhance each other's complementary strengths: the leadership, teamwork, creativity, and social skills of the former, and the speed, scalability, and quantitative capabilities of the latter' (Wilson and Daugherty 2018).

However, with advances in artificial intelligence, machines improve, and might indeed replace humans in their jobs. It is uncertain that enough new jobs at the right skill levels will evolve for everybody, similar to previous shifts in

job markets such as the Industrial Revolution. The demanded skill level might just be too high for all human beings to be able to find a job not yet done by a machine. Or, there just might not be enough jobs left, as more jobs are replaced by machines than are newly created. Massive unemployment would result.

In the short to medium term, regulation could certainly help to avoid mass unemployment, at least for a transitional period. Examples are the requirement for companies to spend a certain amount of their budgets saved via the help of AI on training their workers for higher-skilled jobs; or the restriction of the number of hours worked per day in order to distribute the available work across the entire population. However, in the longer run, if machines replace humans as workers, the idea of a universal basic income will be put back on the table. This would trigger a series of fundamental philosophical but also religious debates: questions such as the purpose of life, how to feel useful and what to strive for, are some issues for which society would have to find answers. Ethics and education will play an important role in order to tackle these societal challenges and questions (Kaplan 2020a).

### *About Democracy and Freedom*

Finally, AI progress could represent nothing less than a danger to peace and democracy (Kaplan 2020b). There are at least two ways in which artificial intelligence might constitute a threat to democracy and its mechanisms, endangering the peaceful coexistence of humans and machines: supervision and manipulation.

Using the example of China, we will provide an illustration as to how far the possibilities of artificial intelligence reach with respect to control and supervision. AI is largely embraced by the Chinese government, which uses it to track and monitor its citizens and inhabitants. For each individual, the Chinese government calculates a so-called ‘social credit score’ based on (big) data coming from various different sources such as health and tax records, social media activity, purchasing behaviour, criminal records and so forth. The system also uses facial recognition and images of the 200 million surveillance cameras mounted across the country for data collection and respective score calculation. Good behaviour such as volunteering at an orphanage leads to higher scores; bad behaviour such as littering leads to lower scores. In order to fulfil the score’s aim, i.e., to encourage good behaviour and citizenship, bad scores result in punishments such as not being eligible for bank loans, not being allowed to fly or not being hired by public agencies (Marr 2019).

In addition to control possibilities, artificial intelligence also allows for manipulation, as we now constantly experience with the dissemination of fake news and disinformation on the various social media platforms (Deighton et al. 2011; Kaplan and Haenlein 2010a; Kaplan 2018). Especially in election campaigns, social media are heavily used to manipulate voters. For example, in the

final three months of the 2016 US presidential election, the top 20 false news items on only one social medium – Facebook – led to more comments, likes and shares than did the 20 most influential news stories from approximately 20 major actors in the news sector together (including such outlets as the *New York Times* and the *Washington Post*; Silverman 2016).

This alone gives enough food for thought regarding the manipulative power of AI-based systems. And yet, the next, bigger thing is just around the corner: deepfakes, which are ‘AI-based technology used to produce or alter audio or video content so that it presents something that did not, in fact, occur’ (Kaplan 2020b). This technology allows inserting words in audio or even video format in an individual’s speech that s/he never actually uttered. Thus, one could make a seemingly authentic video of the Pope stating that monogamy is overrated and that everybody should have open relationships. What this means for future elections and other phenomena is indeed difficult to imagine.

The above two examples clearly show that artificial intelligence potentially leads to issues that do not stop at countries’ borders, with Russia having knowingly been deeply involved in the aforementioned 2016 US presidential election. Regulation that applies to some countries only will most likely be ineffective in governing the coexistence of humans and machines. Intensive international coordination and cooperation in regulation is clearly needed, whenever feasible.

Such international cooperation might be a challenge. While China and the United States are considered as the AI superpowers, they are less known for their implementation of AI regulations (Kaplan 2020a). The development of regulation as well as ethics guidelines falls rather within the expertise of the European Union. The EU, however, has far less influence in the actual development and elaboration of artificial intelligence. Nevertheless, spill-over effects are possible. The EU’s strict General Data Protection Regulation (GDPR), effective since May 2018, applies to any corporation that markets products to EU residents, regardless of its location. Thus, GDPR influences data protection requirements worldwide. As such, the California Consumer Privacy Act (CCPA), which governs the most populous US state’s data protection since January 2020, is recurrently referred to as California’s GDPR. Government regulation is certainly a necessary step. Most likely, whenever society realises the topic’s importance, companies will feel obliged to go into the direction of self-regulation, similarly to the worldwide impact of citizens’ increased commitment and desire for sustainability and a stronger protection of the environment.

### **Conclusion: Only Time Will Tell**

In this chapter, we introduced the concept of artificial intelligence and how it differs from related concepts such as big data, the Internet-of-Things, and machine learning. We also surveyed AI’s history and evolution before discussing the relationship between humans and machines from various angles.

Future research will be needed to address the various challenges with regards to the development of artificial intelligence. Which formal method can be used to test for algorithmic bias? Can we identify simple to use measures to assess bias, similar to the way we assess reliability and validity? What is the best way to bridge (deep) learning and privacy? Should learning be conducted on the user side (with algorithms requiring new data)? Or should data be transferred to a trusted intermediary who performs the analysis on behalf of firms? Do users need to be compensated in one way or another for data or resources provided? Moreover, how can the refusal to share data lead to biases in the data available for learning? Which data sources can and should be used for algorithmic learning? Are there certain types of data that should be 'off-limits'? What role will interdisciplinary AI teams play in establishing coexistence between humans and machines? To mention just a few of the potential future research questions, which, in the light of the unprecedented global COVID-19 pandemic and its acceleration of society's digitalisation, become of vital importance.

At least for the moment, it looks as if AI-driven machines will enhance human work instead of replacing it. This is also the opinion of John Kelly, vice president of IBM, who stated, 'Man and machine working together always beat or make a better decision than a man or a machine independently' (Waytz 2019). Moreover, according to a recent Accenture study, more than 60% of employees believe that AI will have a beneficial impact on their work and jobs (Shook and Knickrehm 2017).

COVID-19 impressively showed that artificial intelligence has played an important role in tackling this unprecedented health crisis on a global level. As such, researchers worldwide made use of AI to efficiently identify potentially infected humans, analyse the virus, test possible treatments and therapies, and more generally to find strategies to fight the pandemic. AJ Venkatakrishnan, e.g., applying AI, discovered that a mutation of the original virus would mimic a protein which the human body uses to regulate its fluid and salt equilibrium (Cha 2020). However, the application of artificial intelligence also showed its connected impact on individuals' daily lives as well as on such questions as data security and privacy. Regulation for the human-machine entanglement is clearly needed.

Furthermore, an example at Mercedes-Benz clearly shows that the replacement of the human workforce is still not as easy as sometimes claimed, and that indeed, currently, human < > machine coexistence is here. Normally, in the automobile manufacturing process, robots and automation are common. However, Mercedes-Benz key accounts increasingly demand more customisation – which the robots were not able to deliver.

Therefore, the German automobile giant decided to replace the fully automated process with 'cobots', or collaborative robots, which are robots designed to physically interact with human beings in a shared workspace. These cobots are controlled by humans, and are to be considered an extension of the human's body, facilitating the carrying and moving of heavy car parts. This form of human < > machine collaboration enables an efficient and productive

customization process, responding in real time to customers' precise choices with regard to leather seats, tyre caps, and so forth.

As in the automotive sector, AI will certainly trigger changes and evolutions in the upcoming years in many sectors. Without a crystal ball, it will be difficult to know where and how the coexistence of humans and machines will evolve. However, it is crystal clear that the business world (and society at large) will need to constantly adapt to advances in AI in order to keep up with the pace (Kaplan and Haenlein 2020), or, to quote Benjamin Franklin: 'When you're finished changing, you're finished.'

## References

- Asimov, I. 1950. Runaround in *I, Robot: The Isaac Asimov Collection*. New York: Doubleday.
- Cha, A. 2020. Artificial Intelligence and Covid-19: Can the Machines Save Us? *Wall Street Journal*, 1 November 2020.
- Conner-Simons, A. 2016. System predicts 85% of cyber-attacks using input from human experts. *MIT News*. Retrieved from: <http://news.mit.edu/2016/ai-system-predicts-85-percent-cyber-attacks-using-input-human-experts-0418>.
- Darrach, B. 1970. Meet Shaky, the first electronic person, *Life Magazine*, 69(21), pp. 58–68.
- Deighton, J., Fader, P., Haenlein, M., Kaplan, A.M., Libei, B. and Muller, E. 2011. Médias Sociaux et Entreprise, une Route Pleine de Défis. *Recherche et Applications en Marketing*, 26(3), 117–124.
- Gandomi, A. and Haider, M. 2015. Beyond the hype: Big Data Concepts, Methods, and Analytics. *International Journal of Information Management*, 35(2), 137–144.
- Haenlein, M. and Kaplan, A. 2019. A Brief History of AI: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*, 61(4), 5–14.
- Haenlein, M. and Kaplan, A. 2021. Artificial Intelligence And Robotics: Shaking Up the Business World and Society at Large. *Journal of Business Research*, 124, 405–407.
- Huang, M.-H. and Rust, R.T. 2018. Artificial Intelligence in Service. *Journal of Service Research*, 21(2), 155–172.
- Kahn, J. 2002. It's alive. *Wired*, 1 March.
- Kaplan, A. 2012. If you Love Something, Let it Go Mobile: Mobile Marketing and Mobile Social Media 4x4. *Business Horizons*, 55(2), 129–139.
- Kaplan, A. 2018. Social Media. In: B. Warf (Ed.), *The Sage Encyclopedia of the Internet*. London: Sage Publications Ltd.
- Kaplan, A. 2020a. Artificial Intelligence: Emphasis on Ethics and Education. *International Journal of Swarm Intelligence and Evolutionary Computation*, 9(3).

- Kaplan, A. 2020b. Artificial Intelligence, Social Media, and Fake News: Is This the End of Democracy? In A.A. Gül, Y.D. Ertürk and P. Elmer (Eds.), *Digital Transformation in Communication and Media Studies*. Istanbul: Istanbul University Press.
- Kaplan, A. 2020c. Retailing and the Ethical Challenges and Dilemmas behind Artificial Intelligence. In E. Pantano (Ed.), *Retail Futures: The Good, the Bad, and the Ugly of the Digital Transformation*. Bingley: Emerald Publishing.
- Kaplan, A. 2021. *Higher Education at the Crossroads of Disruption: The University of the 21st Century, Great Debates in Higher Education*. Bingley: Emerald Publishing.
- Kaplan, A. and Haenlein, M. 2010a. Uitdagingen en kansen rond social media. *Management Executive*, 8(3), 18–19.
- Kaplan, A. and Haenlein, M. 2010b. Users of the world, unite! The challenges and opportunities of social media. *Business Horizons*, 53(1), 59–68.
- Kaplan, A. and Haenlein, M. 2019. Siri, Siri in my hand, who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25.
- Kaplan, A. and Haenlein, M. 2020. Rulers of the world, unite! The challenges and opportunities of artificial intelligence. *Business Horizons*, 63(2).
- Krotov, V. 2017. The Internet of Things and New Business Opportunities, *Business Horizons*, 60(6), 831–841.
- Libai, B., Bart, Y., Gensler, S., Hofacker, C., Kaplan, A., Köttchenheinrich, K. and Kroll, E. 2020. A brave new world? On AI and the management of customer relationships. *Journal of Interactive Marketing*, 51(C), 44–56.
- Marr, B. 2019. Chinese social credit score: Utopian Big Data – Bliss? Or black mirror on steroids? *Forbes*, 21 January.
- McCordick, P. 2004. *Machines Who Think* (2nd ed.). Natick, MA: A. K. Peters, Ltd.
- Mitchell, T. 1997. *Machine Learning*. New York: McGraw Hill.
- Russell, S.J. and Norvig, P. 2016. *Artificial Intelligence: A Modern Approach*. Malaysia: Pearson Education Limited.
- Saarikko, T., Westergren, U. H., and Blomquist, T. 2017. The Internet of Things: Are You Ready for What's Coming? *Business Horizons*, 60(5), 667–676.
- Shook, E. and Knickrehm, M. 2017. *Harnessing Revolution: Creating the Future Workforce*. Accenture Strategy.
- Silverman, C. 2016. This analysis shows how viral fake election news stories outperformed real news on Facebook. *Buzzfeed*, 16 November.
- Waytz, A. 2019. How humans and machines can live and work together. *CNN Business Perspectives*, 1 May.
- Welch, A. 2018. AI better than dermatologists at detecting skin cancer, study finds. *CBS News*, 29 May.
- Wilson, B., Hoffman, J. and Morgenstern, J. 2020. Predictive Inquiry in Object Detection. Working paper.
- Wilson, J. and Daugherty, P. 2018. Collaborative Intelligence: Humans and AI Are Joining Forces. *Harvard Business Review*, 1 July.