

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/334539401>

A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence

Article in *California Management Review* · July 2019

DOI: 10.1177/0008125619864925

CITATIONS

1,839

READS

129,691

2 authors:



Michael Haenlein

ESCP Business School

69 PUBLICATIONS 37,997 CITATIONS

SEE PROFILE



Andreas Kaplan

Kühne Logistics University

98 PUBLICATIONS 31,829 CITATIONS

SEE PROFILE

A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence

California Management Review
1–10

© The Regents of the

University of California 2019

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/0008125619864925

journals.sagepub.com/home/cmrv**Michael Haenlein¹ and Andreas Kaplan²**

SUMMARY

This introduction to this special issue discusses artificial intelligence (AI), commonly defined as “a system’s ability to interpret external data correctly, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation.” It summarizes seven articles published in this special issue that present a wide variety of perspectives on AI, authored by several of the world’s leading experts and specialists in AI. It concludes by offering a comprehensive outlook on the future of AI, drawing on micro-, meso-, and macro-perspectives.

KEYWORDS: artificial intelligence, big data, regulation, strategy, machine-based learning

The world we are living in today feels, in many ways, like a Wonderland similar to the one that the British mathematician Charles Lutwidge Dodgson, better known under the name Lewis Carroll, described in his famous novels. Image recognition, smart speakers, and self-driving cars—all of this is possible due to advances in artificial intelligence (AI), defined as “a system’s ability to interpret external data correctly, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation.”¹ Established as an academic discipline in the 1950s, AI remained an area of relative scientific obscurity and limited practical interest for over half a century. Today, due to the rise of Big Data and improvements in computing power, it has entered the business environment and public conversation.

¹ESCP Europe Business School, Paris, France

²ESCP Europe Business School, Berlin, Germany

AI can be classified into analytical, human-inspired, and humanized AI depending on the types of intelligence it exhibits (cognitive, emotional, and social intelligence) or into Artificial Narrow, General, and Super Intelligence by its evolutionary stage.² What all of these types have in common, however, is that when AI reaches mainstream usage it is frequently no longer considered as such. This phenomenon is described as the AI effect, which occurs when onlookers discount the behavior of an AI program by arguing that it is not real intelligence. As the British science fiction writer Arthur Clarke once said, “Any sufficiently advanced technology is indistinguishable from magic.” Yet when one understands the technology, the magic disappears.

In regular intervals since the 1950s, experts predicted that it will only take a few years until we reach Artificial General Intelligence—systems that show behavior indistinguishable from humans in all aspects and that have cognitive, emotional, and social intelligence. Only time will tell whether this will indeed be the case. But to get a better grasp of what is feasible, one can look at AI from two angles—the road already traveled and what still lies ahead of us. In this editorial, we aim to do just that. We start by looking into the *past* of AI to see how far this area has evolved using the analogy of the four seasons (spring, summer, fall, and winter), then into the *present* to understand which challenges firms face today, and finally into the *future* to help everyone prepare for the challenges ahead of us.

The Past: Four Seasons of AI

AI Spring: The Birth of AI

Although it is difficult to pinpoint, the roots of AI can probably be traced back to the 1940s, specifically 1942, when the American Science Fiction writer Isaac Asimov published his short story *Runaround*. The plot of *Runaround*—a story about a robot developed by the engineers Gregory Powell and Mike Donovan—evolves around the *Three Laws of Robotics*: (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 to it by human beings except where such orders would conflict with the First Law; and (3) a robot must protect its own existence as long as such protection does not conflict with the First or Second Laws. Asimov’s work inspired generations of scientists in the field of robotics, AI, and computer science—among others the American cognitive scientist Marvin Minsky (who later co-founded the MIT AI laboratory).

At roughly the same time, but over 3,000 miles away, the English mathematician Alan Turing worked on much less fictional issues and developed a code breaking machine called *The Bombe* for the British government, with the purpose of deciphering the Enigma code used by the German army in the Second World War. *The Bombe*, which was about 7 by 6 by 2 feet large and had a weight of about a ton, is generally considered the first working electro-mechanical computer. The powerful way in which *The Bombe* was able to break the Enigma code, a task previously impossible to even the best human mathematicians, made Turing wonder

about the intelligence of such machines. In 1950, he published his seminal article “Computing Machinery and Intelligence”³ where he described how to create intelligent machines and in particular how to test their intelligence. This Turing Test is still considered today as a benchmark to identify intelligence of an artificial system: if a human is interacting with another human and a machine and unable to distinguish the machine from the human, then the machine is said to be intelligent.

The word Artificial Intelligence was then officially coined about six years later, when in 1956 Marvin Minsky and John McCarthy (a computer scientist at Stanford) hosted the approximately eight-week-long *Dartmouth Summer Research Project on Artificial Intelligence (DSRPAI)* at Dartmouth College in New Hampshire. This workshop—which marks the beginning of the AI Spring and was funded by the Rockefeller Foundation—reunited those who would later be considered as the founding fathers of AI. Participants included the computer scientist Nathaniel Rochester, who later designed the IBM 701, the first commercial scientific computer, and mathematician Claude Shannon, who founded information theory. The objective of *DSRPAI* was to reunite researchers from various fields in order to create a new research area aimed at building machines able to simulate human intelligence.

AI Summer and Winter: The Ups and Downs of AI

The Dartmouth Conference was followed by a period of nearly two decades that saw significant success in the field of AI. An early example is the famous ELIZA computer program, created between 1964 and 1966 by Joseph Weizenbaum at MIT. ELIZA was a natural language processing tool able to simulate a conversation with a human and one of the first programs capable of attempting to pass the aforementioned Turing Test.⁴ Another success story of the early days of AI was the General Problem Solver program—developed by Nobel Prize winner Herbert Simon and RAND Corporation scientists Cliff Shaw and Allen Newell—that was able to automatically solve certain kind of simple problems, such as the Towers of Hanoi.⁵ As a result of these inspiring success stories, substantial funding was given to AI research, leading to more and more projects. In 1970, Marvin Minsky gave an interview to *Life Magazine* in which he stated that a machine with the general intelligence of an average human being could be developed within three to eight years.

Yet, unfortunately, this was not the case. Only three years later, in 1973, the U.S. Congress started to strongly criticize the high spending on AI research. In the same year, the British mathematician James Lighthill published a report commissioned by the British Science Research Council in which he questioned the optimistic outlook given by AI researchers. Lighthill stated that machines would only ever reach the level of an “experienced amateur” in games such as chess and that common-sense reasoning would always be beyond their abilities. In response, the British government ended support for AI research in all except three universities (Edinburgh, Sussex, and Essex) and the U.S. government soon followed the

British example. This period started the AI Winter. And although the Japanese government began to heavily fund AI research in the 1980s, to which the U.S. DARPA responded by a funding increase as well, no further advances were made in the following years.

AI Fall: The Harvest

One reason for the initial lack of progress in the field of AI and the fact that reality fell back sharply relative to expectations lies in the specific way in which early systems such as ELIZA and the General Problem Solver tried to replicate human intelligence. Specifically, they were all Expert Systems, that is, collections of rules which assume that human intelligence can be formalized and reconstructed in a top-down approach as a series of “if-then” statements.⁶ Expert Systems can perform impressively well in areas that lend themselves to such formalization. For example, IBM’s Deep Blue chess playing program, which in 1997 was able to beat the world champion Gary Kasparov—and in the process proved one of the statements made by James Lighthill nearly 25 earlier wrong—is such an Expert System. Deep Blue was reportedly able to process 200 million possible moves per second and to determine the optimal next move looking 20 moves ahead through the use of a method called tree search.⁷

However, Expert Systems perform poorly in areas that do not lend themselves to such formalization. For example, an Expert System cannot be easily trained to recognize faces or even to distinguish between a picture showing a muffin and one showing a Chihuahua.⁸ For such tasks it is necessary that a system is able to interpret external data correctly, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation—characteristics that define AI.⁹ Since Expert Systems do not possess these characteristics, they are technically speaking not true AI. Statistical methods for achieving true AI have been discussed as early as the 1940s when the Canadian psychologist Donald Hebb developed a theory of learning known as Hebbian Learning that replicates the process of neurons in the human brain.¹⁰ This led to the creation of research on Artificial Neural Networks. Yet, this work stagnated in 1969 when Marvin Minsky and Seymour Papert showed that computers did not have sufficient processing power to handle the work required by such artificial neural networks.¹¹

Artificial neural networks made a comeback in the form of Deep Learning when in 2015 AlphaGo, a program developed by Google, was able to beat the world champion in the board game Go. Go is substantially more complex than chess (e.g., at opening there are 20 possible moves in chess but 361 in Go) and it was long believed that computers would never be able to beat humans in this game. AlphaGo achieved its high performance by using a specific type of artificial neural network called Deep Learning.¹² Today artificial neural networks and Deep Learning form the basis of most applications we know under the label of AI. They are the basis of image recognition algorithms used by Facebook, speech recognition algorithms that fuel smart speakers and self-driving cars. This harvest of the fruits of past statistical advances is the period of AI Fall, which we find ourselves in today.

The Present: California Management Review Special Issue on AI

The discussion above makes it clear that AI will become as much part of everyday life as the Internet or social media did in the past. In doing so, AI will not only impact our personal lives but also fundamentally transform how firms take decisions and interact with their external stakeholders (e.g., employees, customers). The question is less whether AI will play a role in these elements but more which role it will play and more importantly how AI systems and humans can (peacefully) coexist next to each other. Which decisions should rather be taken by AI, which ones by humans, and which ones in collaboration will be an issue all companies need to deal with in today's world and our articles in this special issue provide insights into this from three different angles.

First, these articles look into the relationship between firms and employees or generally the impact of AI on the job market. In their article "Artificial Intelligence in Human Resources Management: Challenges and a Path Forward" Tambe, Cappelli, and Yakubovich analyze how AI changes the HR function in firms. Human resource management is characterized by a high level of complexity (e.g., measurement of employee performance) and relatively rare events (e.g., occurrence of recruiting and dismissals), which have serious consequences for both employees and the firm. These characteristics create challenges in the data-generation stage, the machine-learning stage, and the decision-making stage of AI solutions. The authors analyze those challenges, provide recommendations on when AI or humans should take the lead, and discuss how employees can be expected to react to different strategies.

Another article that addresses this issue is "The Feeling Economy: Managing in the Next Generation of AI" by Huang, Rust, and Maksimovic. This article takes a broader view and analyzes the relative importance of mechanical tasks (e.g., repairing and maintaining equipment), thinking tasks (e.g., processing, analyzing, and interpreting information), and feeling tasks (e.g., communicating with people) for different job categories. Through empirical analysis, these authors show that in the future, human employees will be increasingly occupied with feeling tasks since thinking tasks will be taken over by AI systems in a manner similar to how mechanical tasks have been taken over by machines and robots.

Second, the articles in this special issue analyze how AI changes the internal functioning of firms, specifically group dynamics and organizational decision making. In "Organizational Decision-Making Structures in the Age of AI," Shrestha, Ben-Menahem, and von Krogh develop a framework to explain under which conditions organizational decision making should be fully delegated to AI, hybrid (either AI as an input to human decision making or human decisions as an input to AI systems) or aggregated (in the sense that humans and AI take decisions in parallel with the optimal decision being determined by some form of voting). The question of which option should be preferred depends on the specificity of the decision-making space, the size of the alternative set, and decision-making speed as well as the need for interpretability and replicability.

In a similar spirit, Metcalf, Askay, and Rosenberg present artificial swarm intelligence as a tool to allow humans to make better decisions in “Keeping Humans in the Loop: Pooling Knowledge through Artificial Swarm Intelligence to Improve Business Decision Making.” By taking inspiration from decision making in the animal world (e.g., among flocks of birds or ant colonies), these authors propose a framework to combine explicit and tacit knowledge that suffers less from biases such as herding behavior or the limitations of alternative techniques such as surveys, crowdsourcing, or prediction markets. They show the applicability of their method to sales forecasting and the definition of strategic priorities.

In their article “Demystifying AI: What Digital Transformation Leaders Can Teach You,” Brock and Wangenheim take a broader perspective and investigate to what extent firms are already using AI in their business and how leaders in AI are different from companies lagging behind. Based on a large-scale survey, they identify guidelines of successful AI applications that include a need for data, the requirement to have skilled staff and in-house knowledge, the focus on improving existing business offerings using AI, the importance of having AI embedded in the organization (while, at the same time, engaging with technology partners), and the importance of being agile and having top-management commitment.

Finally, the articles in this special issue look into the interaction between a firm and its customers and specifically the role of AI in marketing. In “Understanding the Role of Artificial Intelligence in Personalized Engagement Marketing,” Kumar, Rajan, Venkatesan, and Lecinski propose how AI can help in the automatic machine-driven selection of products, prices, website content, and advertising messages that fit with an individual customer’s preferences. They discuss in detail how the associated curation of information through personalization changes branding and customer relationship management strategies for firms in both developed and developing economies.

In a similar spirit, Overgoor, Chica, Rand, and Weishampel provide a six-step framework on how AI can support marketing decision making in “Letting the Computers Take Over: Using AI to Solve Marketing Problems.” This framework—which is based on obtaining business and data understanding, data preparation and modeling, as well as evaluation and deployment of solutions—is applied in three case studies to problems many firms face in today’s world: how to design influencer strategies in the context of word-of-mouth programs,¹³ how to select images for digital marketing, and how to prioritize customer service in social media.

The Future: Need for Regulation

Micro-Perspective: Regulation with Respect to Algorithms and Organizations

The fact that in the near future AI systems will increasingly be part of our day-to-day lives raises the question of whether regulation is needed and, if so, in which form. Although AI is in its essence objective and without prejudice, it does not mean that systems based on AI cannot be biased. In fact, due to its very

nature, any bias present in the input data used to train an AI system persists and may even be amplified. Research has, for example, shown that the sensors used in self-driving cars are better in detecting lighter skin tones than darker ones¹⁴ (due to the type of pictures used to train such algorithms) or that decision-support systems used by judges may be racially biased¹⁵ (since they are based on the analysis of past rulings).

Instead of trying to regulate AI itself, the best way to avoid such errors is probably to develop commonly accepted requirements regarding the training and testing of AI algorithms, possibly in combination with some form of warranty, similar to consumer and safety testing protocols used for physical products. This would allow for stable regulation even if the technical aspects of AI systems evolve over time. A related issue is the one of accountability of firms for mistakes of their algorithms or even the need for a moral codex of AI engineers, similar to the one lawyers or doctors are swearing to. What such rules can, however, not avoid is the deliberate hacking of AI systems, the unwanted use of such systems for micro-targeting based on personality traits,¹⁶ or the generation of fake news.¹⁷

What makes matters even more complicated is that Deep Learning, a key technique used by most AI systems, is inherently a black box. While it is straightforward to assess the quality of the *output* generated by such systems (e.g., the share of correctly classified pictures), the *process* used for doing so remains largely opaque. Such opacity can be intentional (e.g., if a corporation wants to keep an algorithm secret), due to technical illiteracy or related to the scale of application (e.g., in cases where a multitude of programmers and methods are involved).¹⁸ While this may be acceptable in some cases, it may be less so in others. For example, few people may care how Facebook identifies who to tag in a given picture. But when AI systems are used to make diagnostic suggestions for skin cancer based on automatic picture analysis,¹⁹ understanding *how* such recommendations have been derived becomes critical.

Meso-Perspective: Regulation with Respect to Employment

In a similar manner as the automation of manufacturing processes has resulted in the loss of blue-collar jobs, the rising use of AI will result in less need for white-collar employees and even high-qualified professional jobs. As mentioned previously, image recognition tools are already outperforming physicians in the detection of skin cancer and in the legal profession e-discovery technologies have reduced the need for large teams of lawyers and paralegals to examine millions of documents.²⁰ Granted, significant shifts in job markets have been observed in the past (e.g., in the context of the Industrial Revolution from 1820-1840), but it is not obvious whether new jobs will necessarily be created in other areas in order to accommodate those employees. This is related to both the number of possible new jobs (which may be much less than the number of jobs lost) and the skill level required.

Interestingly, in a similar way as fiction can be seen as the starting point of AI (remember the *Runaround* short story by Isaac Asimov), it can also be used to get a

glimpse into how a world with more unemployment could look like. The fiction novel *Snow Crash* published by the American Writer Neal Stephenson describes a world in which people spend their physical life in storage units, surrounded by technical equipment, while their actual life takes place in a three-dimensional world called the Metaverse where they appear in the form of three-dimensional avatars. As imaginary as this scenario sounds, recent advancements in virtual reality image processing, combined with the past success of virtual worlds²¹ (and the fact that higher unemployment leads to less disposable income), make alternative forms of entertainment less accessible, and make this scenario far from utopian.

Regulation might again be a way to avoid such an evolution. For example, firms could be required to spend a certain percentage of the money saved through automation into training employees for new jobs that cannot be automated. States may also decide to limit the use of automation. In France, self-service systems used by public administration bodies can only be accessed during regular working hours. Or firms might restrict the number of hours worked per day to distribute the remaining work more evenly across the workforce. All of these may be easier to implement, at least in the short term, than the idea of a Universal Basic Income that is usually proposed as a solution in this case.

Macro-Perspective: Regulation with Respect to Democracy and Peace

All this need for regulation necessarily leads to the question “*Quis custodiet ipsos custodes?*” or “Who will guard the guards themselves?” AI can be used not only by firms or private individuals but also by states themselves. China is currently working on a social credit system that combines surveillance, Big Data, and AI to “allow the trustworthy to roam everywhere under heaven while making it hard for the discredited to take a single step.”²² In an opposite move, San Francisco recently decided to ban facial recognition technology²³ and researchers are working on solutions that act like a virtual invisibility cloak and make people undetectable to automatic surveillance cameras.²⁴

While China and, to a certain extent, the United States try to limit the barriers for firms to use and explore AI, the European Union has taken the opposite direction with the introduction of the General Data Protection Regulation (GDPR) that significantly limits the way in which personal information can be stored and processed. This will by all likelihood result in the fact that the development of AI will be slowed down in the EU compared with other regions, which in turn raises the question how to balance economic growth and personal privacy concerns. In the end, international coordination in regulation will be needed, similar to what has been done regarding issues such as money laundering or weapons trade. The nature of AI makes it unlikely that a localized solution that only affects some countries but not others will be effective in the long run.

Through the Looking Glass

Nobody knows whether AI will allow us to enhance our own intelligence, as Raymond Kurzweil from Google thinks, or whether it will eventually lead us

into World War III, a concern raised by Elon Musk. However, everyone agrees that it will result in unique ethical, legal, and philosophical challenges that will need to be addressed.²⁵ For decades, ethics has dealt with the Trolley Problem, a thought experiment in which an imaginary person needs to choose between inactivity which leads to the death of many and activity which leads to the death of few.²⁶ In a world of self-driving cars, these issues will become actual choices that machines and, by extension, their human programmers will need to make.²⁷ In response, calls for regulation have been numerous, including by major actors such as Mark Zuckerberg.²⁸

But how do we regulate a technology that is constantly evolving by itself—and one that few experts, let alone politicians, fully understand? How do we overcome the challenge of being sufficiently broad to allow for future evolutions in this fast-moving world and sufficiently precise to avoid everything being considered as AI? One solution can be to follow the approach of U.S. Supreme Court Justice Potter Stewart who in 1964 defined obscenity by saying: “I know it when I see it.” This brings us back to the AI effect mentioned earlier, that we now quickly tend to accept as normal was used to be seen as extraordinary. There are today dozens of different apps that allow a user to play chess against her phone. Playing chess against a machine—and losing with near certainty—has become a thing not even worth mentioning. Presumably, Garry Kasparov had an entirely different view on this matter in 1997, just a bit over 20 years ago.

Author Biographies

Michael Haenlein is the Big Data Research Center Chaired Professor and Associate Dean of the Executive PhD Program at the ESCP Europe Business School (email: haenlein@escpeurope.eu).

Andreas Kaplan, Professor and Dean at ESCP Europe Business School Berlin, counts among the Top 50 Business and Management authors worldwide (email: akaplan@escpeurope.eu).

Notes

1. Andreas M. Kaplan and Michael Haenlein, “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 (January/February 2019): 15-25.
2. Ibid.
3. Alan Turing, “Computing Machinery and Intelligence,” *Mind*, LIX/236 (1950): 433-460.
4. For those eager to try ELIZA, see: <https://www.masswerk.at/elizabot/>.
5. The Towers of Hanoi is a mathematical game that consists of three rods and a number of disks of different sizes. The game starts with the disks in one stack in ascending order and consists of moving the entire stack from one rod to another, so that at the end the ascending order is kept intact.
6. For more details, see Kaplan and Haenlein, op. cit.
7. Murray Campbell, A. Joseph Hoane Jr., and Feng-Hsiung Hsu, “Deep Blue,” *Artificial Intelligence*, 134/1-2 (January 2002): 57-83.
8. Matthew Hutson, “How Researchers Are Teaching AI to Learn Like a Child,” *Science*, May 24, 2018, <https://www.sciencemag.org/news/2018/05/how-researchers-are-teaching-ai-learn-child>.
9. Kaplan and Haenlein, op. cit.

10. Donald Olding Hebb, *The Organization of Behavior: A Neuropsychological Theory* (New York, NY: John Wiley, 1949).
11. Marvin Minsky and Seymour A. Papert, *Perceptrons: An Introduction to Computational Geometry* (Cambridge, MA: MIT Press, 1969).
12. David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis, "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, 529 (January 27, 2016): 484-489.
13. Michael Haenlein and Barak Libai, "Seeding, Referral, and Recommendation: Creating Profitable Word-of-Mouth Programs," *California Management Review*, 59/2 (Winter 2017): 68-91.
14. Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern, "Predictive Inequity in Object Detection," working paper, February 21, 2017, <https://arxiv.org/pdf/1902.11097.pdf>.
15. Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner, "Machine Bias: There's Software Used across the Country to Predict Future Criminals. And It's Biased against Blacks," *ProPublica*, May 23, 2016, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.
16. Michal Kosinski, David Stillwell, and Thore Graepel, "Private Traits and Attributes Are Predictable from Digital Records of Human Behavior," *Proceedings of the National Academy of Sciences of the United States of America*, 110/15 (2013): 5802-5805.
17. Supasorn Suwajanakorn, Steven M. Seitz, and Ira Kemelmacher-Shlizerman, "Synthesizing Obama: Learning Lip Sync from Audio," working paper, July 2017, https://grail.cs.washington.edu/projects/AudioToObama/siggraph17_obama.pdf.
18. Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society*, 3/1 (June 2016): 1-12.
19. H. A. Haenssle, C. Fink, R. Schneiderbauer, F. Toberer, T. Buhl, A. Blum, A. Kalloo, A. Ben Hadj Hassen, L. Thomas, A. Enk, and L. Uhlmann, "Man against Machine: Diagnostic Performance of a Deep Learning Convolutional Neural Network for Dermoscopic Melanoma Recognition in Comparison to 58 Dermatologists," *Annals of Oncology*, 29/8 (August 2018): 1836-1842.
20. John Markoff, "Armies of Expensive Lawyers, Replaced by Cheaper Software," *The New York Times*, March 4, 2011, <https://www.nytimes.com/2011/03/05/science/05legal.html>.
21. Andreas M. Kaplan and Michael Haenlein, "The Fairyland of Second Life: About Virtual Social Worlds and How to Use Them," *Business Horizons*, 52/6 (November/December 2009): 563-572.
22. "China Invents the Digital Totalitarian State," *The Economist*, December 17, 2016, <https://www.economist.com/briefing/2016/12/17/china-invents-the-digital-totalitarian-state>.
23. "San Francisco Bans Facial Recognition Technology," *The New York Times*, May 14, 2019, <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html>.
24. Simen Thys, Wiebe van Ranst, and Toon Goedeme, "Fooling Automated Surveillance Cameras: Adversarial Patches to Attack Person Detection," working paper, April 18, 2019, <https://arxiv.org/abs/1904.08653>.
25. Andreas M. Kaplan and Michael Haenlein, "Rulers of the World, Unite! The Challenges and Opportunities of Artificial Intelligence," working paper, 2019.
26. Judith Thomson, "Killing, Letting Die, and The Trolley Problem," *Monist: An International Quarterly Journal of General Philosophical Inquiry*, 59/2 (April 1976): 204-217.
27. Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-Francois Bonnefon, and Iyad Rahwan, "The Moral Machine Experiment," *Nature*, 563 (2018): 59-64.
28. Mark Zuckerberg, "The Internet Needs New Rules. Let's Start in These Four Areas," *The Washington Post*, March 30, 2019, https://www.washingtonpost.com/opinions/mark-zuckerberg-the-internet-needs-new-rules-lets-start-in-these-four-areas/2019/03/29/9e6f0504-521a-11e9-a3f7-78b7525a8d5f_story.html.