40

Will Intelligent Machines Rise Up and Overtake Humanity?

Robert J. Marks II

Advances in artificial intelligence (AI) are jaw-dropping. In 2018, a painting made by AI was auctioned at Christie's for $435,000.[1] AI has beat all human competition in games ranging from GO to Texas Hold'em.[2] China uses AI face recognition technology to monitor its citizenry. An unmanned armed drone was used to assassinate Iranian terrorist Qassem Soleimani.[3] Israel's Harpy missile operates using AI. It flies around a predefined kill zone waiting to be illuminated by radar, and then locates and destroys the source of the radar.[4] All this can be done autonomously without human oversight.[5]

Are these accomplishments of AI a red flag for the future? Will we someday be subservient to advanced AI as science fiction movies depict? In *The Terminator*, AI-based Skynet tries to destroy humanity. *The Matrix* depicts a future where AI exploits humanity while keeping people in a happy state of virtual reality distraction.

Some believe AI will usher in a dystopian future. The time at which the intelligence of computers surpasses that of humans is called the *singularity* by Google's Ray Kurzweil.[6] In 1999, he wrote, "Before the next century is over, human beings will no longer be the most intelligent or capable type of entity on the planet."[7]

Likewise, in his bestselling book *Homo Deus*, Yuval Harari posits that the main products of the twenty-first-century economy "will not be textiles, vehicles, and weapons but bodies, brains, and minds." He continues: "…the way humans have treated animals is a good indicator for how [AI] upgraded humans will treat us."[8]

Other top minds, in the roles of AI Chicken Littles, think likewise:

- Physicist Stephen Hawking warns that the emergence of artificial intelligence could be the "worst event in the history of our civilization."[9]
- Henry Kissinger, former Secretary of State under two US presidents, warns: "Philosophically, intellectually—in every way—human society is unprepared for the rise of artificial intelligence."[10]
- Entrepreneur extraordinaire Elon

> Musk says AI is humanity's "biggest existential threat."[11]

If AI is to become this superintelligent threat, presumably it will need to become more intelligent than we are. To achieve this, AI will need to display human traits like understanding, sentience, and creativity. Not everyone agrees this is possible.

Other AI experts believe in human exceptionalism and claim that computer-based AI will never be able to duplicate the human mind. Noted mathematician and physicist Roger Penrose coauthored with Stephen Hawking the theory of black hole singularities.[12] He contends many human attributes, like creativity, are beyond the reach of AI. After discussing a computer-generated conversation between a computer "therapist" and its patient, he observes, "Though this may give an eerie impression that the computer has some understanding, in fact it has none, and is merely following some fairly simple mechanical rules."[13] Elsewhere Penrose is quoted as saying, "Intelligence cannot be present without understanding. No computer has any awareness of what it does."[14]

Gregory Chirikjian, director of the Johns Hopkins University robotics lab, agrees: "[AI does not display human traits] nor will robots be able to exhibit any form of creativity or sentience."[15] In his book *Hit Refresh,* Microsoft CEO Satya Nadella also agrees: "One of the most coveted human skills is creativity, and this won't change. Machines will enrich and augment our creativity, but the human drive to create will remain central."[16]

AI Limitations

Computer limitations support human exceptionalism.[17] These limitations do not mean AI will never be dangerous. Like electricity, dangers in AI will need mitigating. (We'll address the dangers in AI later.) But AI will never try to take over like in the science fiction movies *The Terminator* and *The Matrix*.

Here is one major reason why: Fundamentally, anything computable must follow a step-by-step procedure written in computer code. Step-by-step procedures are called *algorithms*. Computers can only execute algorithms. There are many problems that are provably nonalgorithmic and therefore beyond the capability of computers.

An example of a nonalgorithmic and therefore noncomputable task is taken from Rice's theorem. Can properties of computer code be determined by an algorithm? In other words, can computer software be written to determine what an arbitrary computer code will or won't do? In many important cases, the answer is no. Rice's theorem says computer code cannot be written to examine an arbitrary computer program to determine whether the program at some time will print the number 3. If the first line in the software says "PRINT 3," then the software can easily be identified to print 3. The key to Rice's theorem is that the examining computer code must work for *all* possible programs. This is provably not possible.

A special case of Rice's theorem is the Turing halting problem.[18] It is not algorithmically possible to determine whether an arbitrary computer program will stop or run forever. No computer program can ever be written to determine whether another arbitrary computer program runs forever or halts.

Another nonalgorithmic operation is compression. Given a computer file of an arbitrarily large size, determining how much the file can be compressed is not algorithmic. The smallest a file can be compressed is called its *Kolmogorov complexity*. Above a certain file

size, the Kolmogorov complexity cannot be computed. It is nonalgorithmic.[19]

This leads to an important question: Are there human traits that are nonalgorithmic? If so, these traits will never be computable. If this is true, the human mind will never be able to be uploaded to a computer. Computers can only handle the algorithmic. If you are stripped of your nonalgorithmic properties, you will be either very boring or, more accurately, not you.

Noncomputable human traits include sentience, understanding, and creativity. Each will now be addressed in more detail.

Sentience

A component of sentience is qualia. *Qualia* is an experience from the senses, including taste, smell, and touch.

Let's do a thought experiment. You bite into a segment from an orange. As you bite, the skin on the segment bursts and juice from the bite covers your tongue. You taste a sweet orange flavor as you chew and swallow.

You are now assigned the job of explaining your experience to a man with no sense of taste or smell since birth. How can this be done? You can provide explanations. You can present the man with the chemical components of the segment. He can understand the physics of chewing and the biology of the taste buds. But the true experience of biting and tasting of the juice exploding from the orange's juice vesicles is not possible to communicate to the sense-deprived man. Qualia is beyond description to those without a shared experience.

If the experience of biting a segment from an orange cannot be explained to a man without the senses of taste or smell, how can we expect to duplicate the qualia experience in a computer using computer code? If the true experience of biting into a segment from an orange can't be explained, it is nonalgorithmic and therefore noncomputable.

Many computer operations, including the Turing halting problem and computation of the Kolmogorov complexity, cannot yet be solved by humans. Neither can describing the true experience of biting into an orange segment.

Duplicating qualia is beyond the capability of AI.

Understanding

Philosopher John Searle illustrated that computers do not understand what they do with his example of the "Chinese room."

Searle imagined himself in a room with many file cabinets. A question written in Chinese is slipped through a slot on the door. Searle does not read Chinese. But the file cabinets contain billions of easily searchable questions along with answers written in Chinese. Searle searches through the file cabinets until he finds a match to the question being posed. The file cabinet also contains the answer to the question. When he finds a match, Searle copies the answer in Chinese. He walks to the door and slips the answer through the slot to whoever is on the outside.

From the outside, it appears the occupant of the Chinese room is not only fluent in Chinese but is able understand questions. Not so. The occupant of the Chinese room has no understanding of the Chinese language. In generating the answer to the submitted question, he is simply following the algorithm of pattern matching.

In 2011, IBM's Watson beat champion contestants in the quiz show Jeopardy. Processing natural language, Watson answered queries faster than the human contestants. Watson was the equivalent of a large Chinese room. Instead of file cabinets, Watson had access to all of Wikipedia and then some.

Like Searle in the Chinese room, Watson had no understanding of the meaning of the queries it fielded. It was following an algorithm written by computer programmers.

AI does not understand what it is doing. Understanding is nonalgorithmic and therefore noncomputable.

Creativity

AI has produced some astonishing results. AlphaGo software developed by DeepMind beat the world champion in the complex game of Go. A computer program dubbed Pluribus has beat poker pros in the game of no-limit Texas Hold'em poker.[20] OpenAI's program GPT-3 generates short bursts of astonishingly coherent prose given only a few prompts.[21] Don't these and other computer programs display creativity?

The answer depends on your dictionary. To answer the creativity question, the term *creativity* must first be defined.

The Lovelace test is a simple and easily understood definition of AI creativity offered by Selmer Bringsjord. Named after Ada Lovelace, the first computer programmer, the Lovelace test asks a simple question to determine software creativity: Did the AI in question perform a task that is beyond the explanation of the person who wrote the AI code (or someone with comparable expertise)? If so, the Lovelace test has been passed and AI has been shown to be creative.

Surprise results don't count. Computer programs generate surprising results all the time. They can be a result of faulty computing or generating results outside expectations.

In beating the world champion at Go, the AI program AlphaGo made an unusual move that surprised many, including the Go champion Lee Sedol. Was this creative? Not according to the Lovelace test. AlphaGo was written to play Go. And that's what it did. If AlphaGo could provide an explanation of the game of Go when queried to assist you with your taxes, it would be creative. But AlphaGo was not programmed to explain the game of Go or to fill out your taxes. It was written to play the game of Go and nothing else.

To achieve superintelligence based on AI writing better and better AI software, creativity is required. If computer programs cannot be creative as defined by the Lovelace test, they cannot write better programs beyond the explanations of the original programmer. For this reason alone, superintelligence is not achievable.

To date, no one has successfully claimed a computer program has passed the Lovelace test.

Arguing Outside Your Silo

The Future of AI

Despite the evidence just presented that AI will never display creativity, understanding, or sentience, there are a lot of people who believe that superintelligence will be achieved. Many see a dystopian future where AI will ruin our society. We'll call them AI dystopians. As we have quoted, Bill Gates and Stephen Hawking are in this highly populated camp. But consensus should not be taken as evidence in the status of science or its future. Michael Crichton, author of classic science fiction books like *Jurassic Park*, said as much in a talk given at CalTech:

> There is no such thing as consensus science. If it's consensus, it isn't science. If it's science, it isn't consensus. Period.[22]

But the question under discussion is not about the current state of AI, it is over the future of AI. The limitations of AI rest largely on computer science, so this foundational

knowledge is essential in predicting what AI will accomplish.

Ideology is a contributing factor to opinions about AI of the future. According to technology prophet George Gilder, those who believe AI will exceed the intelligence of humans suffer from a "materialist superstition" that assumes the mind is a meat computer and can therefore be replicated by AI.[23] Ignorance of the computer science thus far presented is also a factor. Those who code and use computer software for a living are often unaware of the limitations placed on AI by established computer science theory. This background is not required to be an excellent coder.

Algorithmic information theory (AIT) is the study of algorithms in computer science.[24] AIT addresses what computer codes can and can't do.[25] Those without a background in AIT may not have solid ground for appreciating or defending the limitations of AI. AIT is a foundational tool for establishing AI limitations, and those who are trained to code for a living may not be familiar with AIT. Great coders do not need to know about AIT.

We have quoted Microsoft CEO Satya Nadella, Johns Hopkins University robotics scientist Gregory Chirikjian, and Oxford mathematician and physicist Sir Roger Penrose. All believe computers will never be creative. All have backgrounds in computer science and mathematics. But so do some of the AI dystopians.

Often, celebrity AI dystopians speak outside of their silo of expertise.

Confusion of Expertise

Ever wonder why actor Kevin Costner testified in front of Congress on the topic of oil spills, as did Ben Affleck on a children's project, and quizmaster Bob Barker on the Captive Elephant Accident Prevention Act?[26] Many people, including some in Congress and the media, apparently equate celebrity in one area to across-the-board genius in everything.

News commentator Laura Ingraham disagrees; she has told clueless politically pontificating music celebrities to *Shut Up and Sing*.[27] Similarly, innovative comedian Ricky Gervais informed celebrity loudmouths, "You know nothing about the real world. Most of you spent less time in school than Greta Thunberg."[28]

One statistic to test the accuracy of prophetic statements is examination of credentials. Kevin Costner is no expert in oil spills and quizmaster Bob Barker has no credentials in elephant accident prevention. They are listened to because they are celebrities. The same analysis can be applied to business and science celebrities who make predictions about AI. Celebrity scientists and business tycoons are often not experts in AI.

In critiquing a person's background, care must be taken. In a debate, citing the lack of credentials of your opponent is called the *genetic fallacy*. The opponent's background is made the issue rather than the topic of debate. That is not the purpose here. The reasons underlying AI limitation have been established earlier in this chapter. The key debate issues have thus been addressed. We are now asking whether the AI dystopians are aware of the simple AIT-related limitations we have discussed.

Henry Kissinger, previously quoted, is alarmed about the impact of AI if not bridled. Kissinger is a gifted politician, diplomat, and geopolitical consultant. But his knowledge of AI apparently comes from reading the writings of others—many of whom adhere to materialist superstition. I suspect Kissinger is unaware of the deeper issues of computer science.

Elon Musk, whom we also quoted previously, is an AI dystopian. He has said, "With artificial intelligence we are summoning the demon."[29] Musk received undergraduate degrees in economics and physics from the University of Pennsylvania and then attended Stanford for a graduate degree in physics but dropped out. He is a hugely successful entrepreneur, having foundational roles in businesses such as PayPal, SpaceX, Tesla Motors, and OpenAI. During 2020 and 2021, Musk floated between being the seventh richest, the second richest, or even the richest man in the world with a net worth reportedly over $150 billion.[30] But does he have the computer science foundation to appreciate the limits of AI? Does expertise in business translate to expertise in AI?

Facebook CEO Mark Zuckerberg doesn't think so. He called Elon Musk's doomsday AI predictions "pretty irresponsible."[31] Business guru George Gilder questions Musk's opinions outside of the silo of business. He notes, "I think Elon Musk is a tremendous entrepreneur, yet he's a quite retarded thinker."[32]

Ray Kurzweil, a successful inventor and businessman, received a BS degree from MIT in computer science. Kurzweil is an avid supporter of superintelligence and has written books on the topic. Superintelligence assumes AI software will write better software that writes better software and on and on. According to Kurzweil, AI will soon be smarter than humans. But if AI writes better AI software not anticipated by the original programmer, then the AI is being creative. It would pass the Lovelace test. Kurzweil, with an undergraduate degree in computer science and vast experience in business and technical innovation, believes AI can be creative. Carefully defined by the Lovelace test, however, computer software cannot be creative.

Kurzweil once worked for Bill Gates at Microsoft. Does Bill Gates—another AI dystopian—have a background in computer science and AIT? Apparently not. As an undergraduate, Gates dropped out of Harvard University to pursue the founding of Microsoft. He is a talented entrepreneur whose success does not require deep studies in computer science. Gates was a knowledgeable programmer with early computer hardware. But much of his success came from his business instincts and his team of lawyers. Gates's father was a named partner in the Seattle law firm Preston Gates & Ellis. When I consulted for Microsoft, my first meeting was with Microsoft lawyers who told me, in no uncertain terms, my legal responsibilities. As expected, I was told Microsoft would own 100 percent of any intellectual property I created. I was instructed not to look at any patents associated with my assigned duty. This was new to me. If Microsoft was sued relating to my work and lost, punitive damages would kick in if I had looked at patents. Doing so could be construed as plagiarizing intellectual property. If I didn't look at patents, my contribution would be ruled a coincidental discovery by the courts and only monetary damages could be recovered.

In my experience, consulting typically requires the signing of documents like nondisclosures and specifying who owns what intellectual property. But Microsoft is my only consulting experience that started with a nose-to-nose meeting with a gaggle of lawyers.

Microsoft's success is due in large part to smart business dealings and not creativity. Their first historical coup was acquisition of MS-DOS. Microsoft did not write DOS. It was purchased by Microsoft in the early 1980s. Because of the rise in popularity of the

IBM PC, MS-DOS became a cash cow for Microsoft.

Microsoft continued to expand, not necessarily by innovation but by acquisition, copying technology, and court battles. Flagship Microsoft software tools that were purchased or copied from other companies and not innovated by Microsoft include Windows, Word, PowerPoint, Excel, Internet Explorer, and Bing. Business practice, not innovation, is the secret of Microsoft's success.

Bill Gates must be celebrated as a gifted entrepreneur, businessman, and philanthropist. But his background in computer science, other than coding and its application, must be questioned. AI dystopian Gates opined, "I do think we do have to worry about [AI] but I don't think it's inherent that as we create super intelligence that it will necessarily have the same goals in mind that we do."[33]

So like Ray Kurzweil, Bill Gates believes in AI creativity that leads to superintelligence.

Great scientists risk similar overreach in expertise. Consider, for example, Stephen Hawking's fears of AI. Hawking, whose celebrity appearances include *Star Trek: The Next Generation*, is a genius in cosmology. With fellow genius Roger Penrose, he formulated the Penrose-Hawking singularity theorems, applying general relativity so as to better understand black holes. But artificial intelligence disturbed Hawking. He told BBC:

> The development of full artificial intelligence could spell the end of the human race…It [AI] would take off on its own, and re-design itself at an ever-increasing rate…Humans, who are limited by slow biological evolution, couldn't compete, and would be superseded.[34]

Like Gates and Kurzweil, Hawking bought into the idea of superintelligence. Despite his gifts in cosmology, Hawking was outside of his silo of expertise. He assumed that AI is creative.

Hawking's concern was all the more curious in light of his earlier abandonment of an ultimate Theory of Everything that would unify the physics of the universe in a nicely wrapped interconnected set of equations. He changed his mind about the viability of that project on account of Kurt Gödel's landmark theorems on incompleteness and inconsistency. No matter how much physicists discovered, he concluded, there would still be more to learn.[35]

Remarkably, the distance between Gödel's work and the limitations of computer creativity is not far. Indeed, Gödel's theorems form the foundation for AIT. Alan Turing, the father of modern computer science, built on Gödel's thesis, showing that some problems are nonalgorithmic and therefore cannot be captured by computer code. The nonalgorithmic nature of the Turing halting problem discovered by Turing is a manifestation of Gödel's work.

Roger Penrose, Hawking's coauthor on black hole physics, recognized this connection and wrote about it. His book *The Emperor's New Mind* (1989) wonderfully links Gödel to Turing and establishes the noncomputable nature of creativity. Penrose posits that the noncomputable characteristics of human thinking spring from quantum effects in microtubules in the brain. The quantum collapse of a wave function to a deterministic state, reasoned Penrose, is nonalgorithmic and thus might be the source of noncomputable creativity.

Hawking seems not to have considered Penrose's work when he offered his scary prediction about a dystopian AI future. One

would not say to a man of his intellectual stature, "Shut up and do physics!" We can say, however, that Hawking had the intellect to learn about Penrose's insights into computers and AI via Gödel and thereby understand the limits of computers.

Yet apparently he didn't.

AI Dangers

The danger of AI lies not in its potential ability to become conscious and take over the world, but in its incomplete vetting prior to use. AI, like any system, will have its unexpected consequences. Contingencies can increase exponentially as the complexity is increased linearly.

Undesirable and unexpected contingencies have already been manifested in the deployment of AI systems. There are multiple examples of unexpected contingencies arising from complex systems ranking from the simply curious to the very serious:

1. A deep convolutional neural network was trained to detect wolves. After the trained neural network incorrectly classified a husky dog as a wolf, the programmers did some forensics and discovered there was undesirable bias in the training data. The pictures of wolves all contained snow. The picture of the misclassified dog also contained snow. In training, the neural network had learned the presence and absence of snow. The features of the animals were not considered in the classification problem.

2. An inconvenience for self-driving cars is the false classification of objects like plastic bags. A stationary plastic bag can be categorized as a large rock[36] while a wind-blown plastic bag may be mistaken for a deer.[37] These are unintended contingencies of the self-driving car's software.

3. A more serious problem with self-driving cars is fatalities. In 2018, an Uber self-driving car in Tempe, Arizona struck and killed pedestrian Elaine Herzberg. Steven Shladover, a UC Berkeley research engineer, noted, "I think the sensors on the vehicles should have seen the pedestrian well in advance."[38] The death was a tragic example of an unintended contingency of a complex AI system. Unintended contingencies remain a major obstacle in the development of general (level 5) self-driving cars. Some developers, believing the problem in insurmountable, have given up.[39]

4. During the height of the Cold War, the Soviets deployed a satellite early-warning system called Oko to watch for incoming missiles fired from the United States. On September 26, 1983, Oko detected incoming missiles. At a military base outside of Moscow, sirens blared and the Soviet brass was told by Oko to launch a thermonuclear counterstrike. Doing so would result in millions being killed. The officer in charge, Lieutenant Colonel Stanislav Petrov, felt something was fishy. After informing his superiors of his hunch that Oko was not operating correctly, Petrov did not obey the order. Upon further investigation, Oko was found to have mistakenly interpreted the effects of the sun reflecting off of clouds.[40] There was no US missile

attack. Petrov's skepticism of Oko's alarm likely saved millions of lives.

These examples of unintended contingencies deal with systems of broad complexity. Narrow AI systems are typically more error-free. Examples of a narrow AI system are anti-radiation missiles like the previously mentioned Harpy missile from Israel. The missile is launched and flies about (loiters) over a predefined kill zone. The missile can operate autonomously without human oversight. If fuel gets low, the missile returns home. Alternately, if illuminated by radar, the anti-radiation missile zeroes in on the location of the radar's source. The missile follows the radar beam and destroys the radar installation.[41] Whether or not one agrees with the mission of such a system, the anti-radiation missile is an example of relatively narrow AI that has historically worked without flaw. There are few if any unforeseen contingencies in anti-radiation missiles that would distract from their duties.

Self-driving cars require a more complex AI system. For tightly connected AI systems, the number of contingencies and therefore the number of unexpected consequences increase exponentially as the complexity increases linearly.

The dangers of AI can be mitigated through proper design practices:

1. *Domain expertise.* AI software should be developed by those with experiential knowledge of the problem being solved. Experts will better identify undesired contingencies during development of the AI software.
2. *Testing.* AI systems must be tested under different conditions and in different environments.
3. *Disjunctive design.* AI consisting of conjoined narrow AI systems is more easily testable than tightly designed systems. Smaller systems are easier to both conceive and test.

Exploding contingency count should give pause to anyone designing complex AI systems. Even if the development of superintelligence were possible, the required complexity of such systems would present programmers with the overwhelming problem of eliminating unexpected and undesirable performance.

Takeaways

The nonalgorithmic capacity of the human mind remains beyond the reach of computers and AI.[42] But, as with all new technologies like electricity, care must be taken in the development of AI to assure its safe and proper use. Frayed electrical wires still burn down houses and downed electric lines still electrocute. But the advantages of electricity far outweigh the dangers. The negative consequences have been mitigated by legislation and best practices. The hope is that AI dangers can be likewise contained.

AI will never be sentient, creative, nor have understanding, and it will never have dominion over humans because of its own initiative.[43] If AI becomes dangerous, it will be the fault of humans who develop and use AI. Following good ethics will likewise never be the responsibility of the AI itself, but will always belong to those who write and test AI systems and to the end users (humans).

The limitations of AI are further evidence of human exceptionalism.[44] We are "fearfully and wonderfully made."[45]

Chapter 40—Will Intelligent Machines Rise Up and Overtake Humanity?

1. "Is Artificial Intelligence Set to Become Art's Next Medium?" *Christie's* (December 12, 2018), https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx (accessed August 24, 2020).
2. "Google AI defeats human Go champion," BBC News (May 25, 2017), https://www.bbc.com/news/technology-40042581 (accessed August 24, 2020); Carnegie Mellon University, "AI beats professionals in six-player poker," ScienceDaily (July 11, 2019), https://www.sciencedaily.com/releases/2019/07/190711141343.htm (accessed August 24, 2020).
3. Ken Dilanian and Courtney Kube "Airport informants, overhead drones: How the U.S. killed Soleimani," *NBC News* (January 10, 2020), https://www.nbcnews.com/news/mideast/airport-informants-overhead-drones-how-u-s-killed-soleimani-n1113726 (accessed August 24, 2020).
4. Charlie Gao, "The Ultimate Weapon of War No One Is Talking About," *The National Interest* (January 25, 2019), https://nationalinterest.org/blog/buzz/ultimate-weapon-war-no-one-talking-about-42497 (accessed August 24, 2020).

5. Robert J. Marks, *The Case for Killer Robots* (Seattle, WA: Discovery Institute Press, 2020).
6. Ray Kurzweil, *The Singularity Is Near: When Humans Transcend Biology* (London, UK: Penguin, 2005).
7. Ray Kurzweil, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence* (New York: Penguin, 2000).
8. Yuval Noah Harari. *Homo Deus: A Brief History of Tomorrow* (New York: Random House, 2016).
9. Arjun Kharpal, "Stephen Hawking says A.I. could be 'worst event in the history of our civilization,'" *CNBC* (November 6, 2017), https://www.cnbc.com/2017/11/06/stephen-hawking-ai-could-be-worst-event-in-civilization.html (accessed August 24, 2020).
10. Henry Kissinger, "How the Enlightenment Ends," *The Atlantic* (August 30. 2019), https://www.theatlantic.com/magazine/archive/2018/06/henry-kissinger-ai-could-mean-the-end-of-human-history/559124/ (accessed August 24, 2020).
11. Samuel Gibbs, "Elon Musk: artificial intelligence is our biggest existential threat," *The Guardian* (October 27, 2014), https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat (accessed August 24, 2020).
12. Stephen Hawking, Stephen William, and Roger Penrose, "The singularities of gravitational collapse and cosmology," *Proceedings of the Royal Society of London A: Mathematical and Physical Sciences* 314 (1970), 529-548.
13. Roger Penrose, *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics* (Oxford, UK: Oxford University Press, 1989), 16.
14. "Quotations Roger Penrose," https://www-history.mcs.st-andrews.ac.uk/Quotations/Penrose.html (accessed August 24, 2020).
15. Gregory Chirikjian, "Help Wanted: For the Cognitive Era," *JHU Engineering Magazine* (May 19, 2017), https://engineering.jhu.edu/magazine/2017/05/help-wanted-for-the-cognitive-era/ (accessed August 24, 2020).
16. Satya Nadella, Greg Shaw, and Jill Tracie Nichols, *Hit Refresh: The Quest to Rediscover Microsoft's Soul and Imagine a Better Future for Everyone* (New York: Harper Business), 207.
17. Some of the material in this chapter is borrowed from the author's works posted on MindMatters.AI, including "Advice to Physicists: 'Shut Up and Do Physics'" (May 25, 2020), https://mindmatters.ai/2020/05/advice-to-physicists-shut-up-and-do-physics/ (August 24, 2020) and the paper "Linear Complexity Can Generate an Exponential Explosion of Contingencies: Implications for General Artificial Intelligence," currently under review for publication with two coauthors.
18. Robert J. Marks II, William A. Dembski, and Winston Ewert, *Introduction to Evolutionary Informatics* (Singapore: World Scientific, 2017).
19. Thomas M. Cover and Joy A. Thomas, *Elements of information theory* (Hoboken, NJ: John Wiley & Sons, 1999).
20. Meilan Solly, "This Poker-Playing A.I. Knows When to Hold 'Em and When to Fold 'Em," Smithsonian (July 19, 2019), https://www.smithsonianmag.com/smart-news/poker-playing-ai-knows-when-hold-em-when-fold-em-180972643/ (accessed August 24, 2020).
21. Rob Toews, "GPT-3 Is Amazing—And Overhyped," *Forbes* (July 19, 2020), https://www.forbes.com/sites/robtoews/2020/07/19/gpt-3-is-amazingand-overhyped/ (accessed August 24, 2020).
22. Quoted in Jorge R. Barrio, "Consensus Science and the Peer Review," *Molecular Imaging and Biology* 11 (2009), 293.
23. George Gilder, "The Materialist Superstition," *The Intercollegiate Review* (Spring 1996), 6-14.
24. Gregory J. Chaitin, "Algorithmic information theory," *IBM Journal of Research and Development* 21 (1977), 350-359.
25. Peter D. Grünwald and Paul MB Vitányi, "Algorithmic information theory," *Handbook of the Philosophy of Information* (2008), 281-320.
26. Debra Bell, "Celebrities Who Have Testified to Congress," *US News & World Report* (September 24, 2010), https://www.usnews.com/news/slideshows/celebrities-who-have-testified-to-congress (accessed August 25, 2020).
27. Laura Ingraham, *Shut Up and Sing: How Elites from Hollywood, Politics, and the Media Are Subverting America* (Washington, DC: Regnery, 2003).
28. Keiran Southern, "'Most of you spent less time in school than Greta Thunberg'—Ricky Gervais's best and most shocking jokes from the Golden Globes," *Independent.ie* (January 6, 2020), https://www.independent.

ie/entertainment/movies/most-of-you-spent-less-time-in-school-than-greta-thunberg-ricky-gervaiss-best-and-most-shocking-jokes-from-the-golden-globes-38837007.html (accessed August 25, 2020).

29. Matt McFarland, "Elon Musk: 'With artificial intelligence we are summoning the demon,'" *The Washington Post* (October 24, 2014), https://www.washingtonpost.com/news/innovations/wp/2014/10/24/elon-musk-with-artificial-intelligence-we-are-summoning-the-demon/ (accessed August 25, 2020).
30. Tom Maloney, "Elon Musk Soars Past Warren Buffett on Billionaires Ranking," *Bloomberg News* (July 10, 2020), https://www.bloomberg.com/news/articles/2020-07-10/elon-musk-rockets-past-warren-buffett-on-billionaires-ranking (accessed March 16, 2021); Robert Frank, "Elon Musk is now the richest person in the world, passing Jeff Bezos," CNBC (January 7, 2021), https://www.cnbc.com/2021/01/07/elon-musk-is-now-the-richest-person-in-the-world-passing-jeff-bezos-.html (accessed March 16, 2021); Annie Palmer, "Jeff Bezos overtakes Elon Musk to reclaim spot as world's richest person," CNBC (February 16, 2021), https://www.cnbc.com/2021/02/16/jeff-bezos-reclaims-spot-as-worlds-richest-person-from-elon-musk.html (accessed March 16, 2021).
31. Catherine Clifford, "Facebook CEO Mark Zuckerberg: Elon Musk's doomsday AI predictions are 'pretty irresponsible,'" *CNBC* (July 24, 2017), https://www.cnbc.com/2017/07/24/mark-zuckerberg-elon-musks-doomsday-ai-predictions-are-irresponsible.html (accessed August 25, 2020).
32. Rich Karlgaard, "Why Technology Prophet George Gilder Predicts Big Tech's Disruption" *Forbes* (February 8, 2018), https://www.forbes.com/sites/richkarlgaard/2018/02/09/why-technology-prophet-george-gilder-predicts-big-techs-disruption/ (accessed August 25, 2020).
33. "Bill Gates: I think we do need to worry about artificial intelligence," *Fox Business* YouTube channel (January 22, 2016), https://www.youtube.com/watch?v=EmfrMKLwr3k (accessed August 25, 2020).
34. Rory Cellan-Jones, "Stephen Hawking warns artificial intelligence could end mankind," *BBC News* (December 2, 2014), bbc.com/news/technology-302290540.
35. Stephen Hawking, "Gödel and the End of Physics," TAMU lecture (2017), http://yclept.ucdavis.edu/course/215c.S17/TEX/GodelAndEndOfPhysics.pdf (accessed August 27, 2020).
36. Linghe Kong, Muhammad Khurram Khan, Fan Wu, Guihai Chen, and Peng Zeng, "Millimeter-wave wireless communications for IoT-cloud supported autonomous vehicles: Overview, design, and challenges," *IEEE Communications* 55 (2017), 62-68.
37. Neil McBride, "The ethics of driverless cars," *ACM SIGCAS Computers and Society* 45 (2016), 179-184.
38. Aarian Marshall, "The Uber Crash Won't Be the Last Shocking Self-Driving Death," *Wired* (March 31, 2018), https://www.wired.com/story/uber-self-driving-crash-explanation-lidar-sensors/ (accessed August 28, 2020).
39. Brendan Dixon, "Star Self-Driving Truck Firm Shuts; AI Not Safe Enough Soon Enough," *Mind Matters News* (March 24, 2020), https://mindmatters.ai/2020/03/star-self-driving-truck-firm-shuts-ai-not-safe-enough-soon-enough/ (accessed August 28, 2020).
40. Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: Norton, 2018).
41. Marks, *The Case for Killer Robots*.
42. Roger Penrose identified quantum collapse as nonalgorithmic and conjectured this might be the source of humans' nonalgorithmic abilities. There has been no research performed yet, and Penrose's highly speculative hypothesis remains untested.
43. Psalm 8:6; see also Genesis 1.
44. See Genesis 1:27 and 1 Corinthians 2:16.
45. Psalm 139:14.