The Future Isn't What It Used to Be:[1]
Artificial Intelligence Meets Natural Stupidity[2]

Edited version of a presentation given at the
Second Interpreter Science and Mormonism Symposium: Body, Brain, Mind, and Spirit
Utah State University, Orem, Utah
12 March 2016

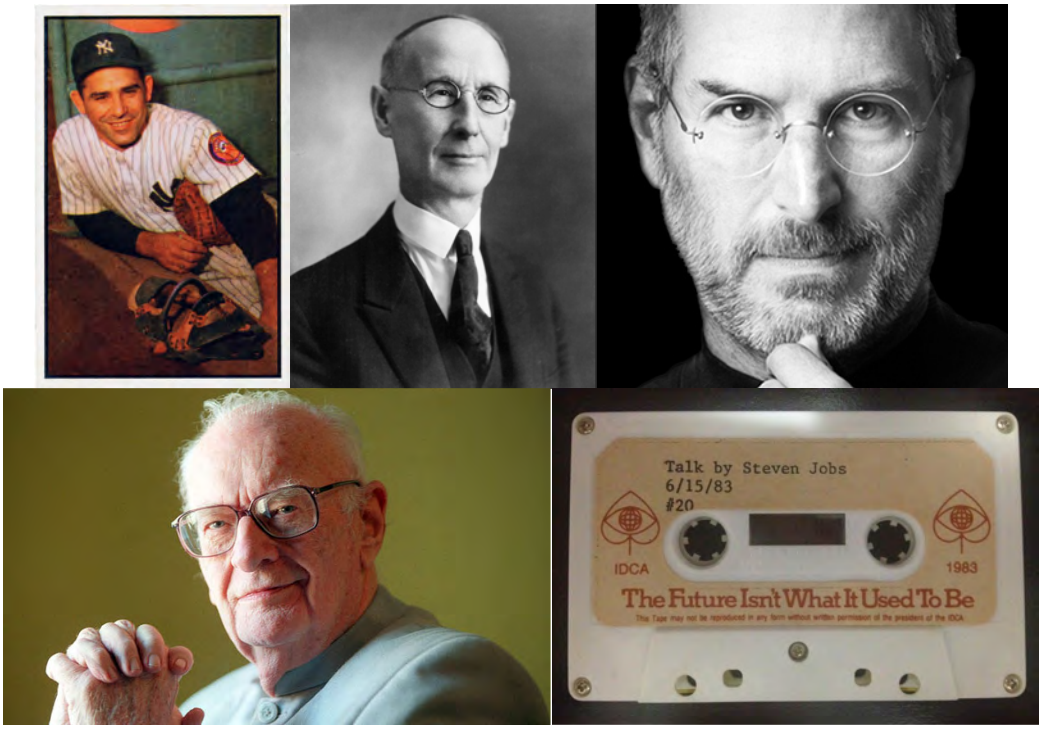

3

The saying "The future isn't what it used to be," is often attributed to Yogi Berra,[4] although he admitted (as one might also conclude about Elder J. Golden Kimball[5]), that he didn't really say everything he said.[6] The same aphorism was used by scientist and science fiction author Sir Arthur C. Clarke[7] and later by Apple Computer co-founder Steve Jobs[8] as a preface to their optimistic extrapolations about the future of technology. Although I would agree with Clarke and Jobs that a bright future for technology lies ahead, that is not the kind of talk I intend to give today.

More significant than the transformations that technology works upon us directly,[9] more potent than any hallucinogenic drug on our thoughts and senses are the mind-altering effects of our changing conceptions about the future itself. Arthur C. Clarke observed: "Until a century ago nobody was very interested in the future for the simple reason that, apart from natural catastrophes and wars, the future was going to be the same as the past. A man knew that the pattern of his life would be the same as his great-grandfather's, as far back as anyone could remember. Well, now we know differently."[10]

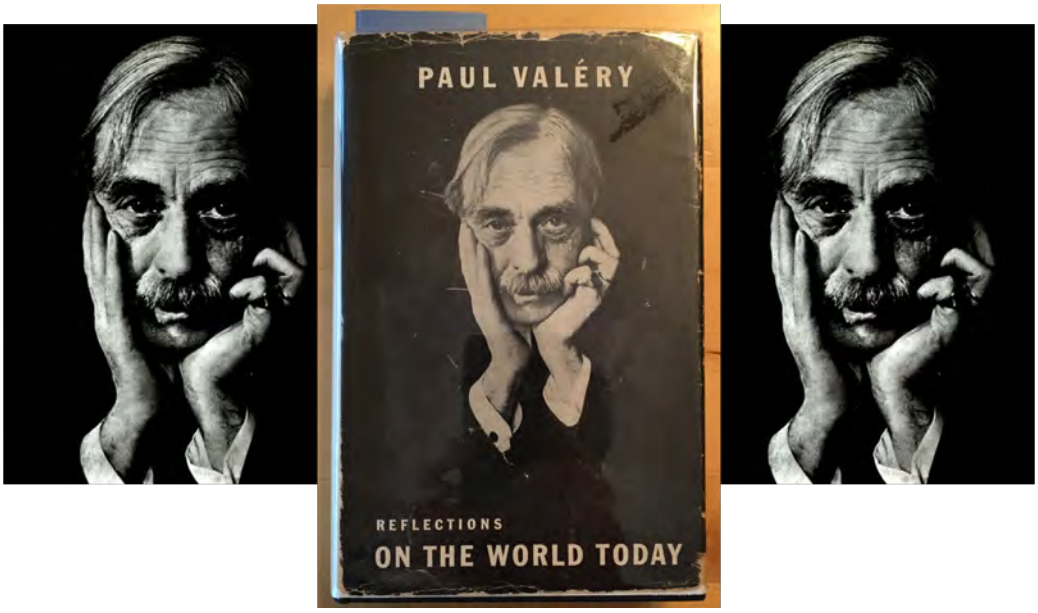

11

The brilliant, problematic French poet and essayist Paul Valéry (1871-1945)[12] explained why "the future isn't what it used to be" in 1937, long before any of the others I cited before.[13] Wrote Valéry:[14]

> The future, like everything else, is no longer quite what it used to be. … We used to consider the unknown future as a simple combination of already known things, and the new was analyzed according to its unoriginal elements. But that is ended. … [T]he rules of the game are changed at every throw. No calculation of probabilities is possible. … Why? Because the … modern world is assuming the shape of man's mind. Man has sought in nature all the means and powers that are necessary to make the things around him as unstable, volatile, and mobile as himself, as admirable, as absurd, as disconcerting and prodigious as his own mind. … If … we imprint the form of our mind on the human world, the world becomes all the more unforeseeable and assumes the mind's [own] disorder.

"Happily," concludes Valéry, "these prophecies are idle. I am busy doing exactly what I explained the vanity of a few moments ago. I am looking ahead, therefore I am wrong."[15]

Since the *far* future is impossible to predict with any confidence, my burden today will be to share some candid observations about the present and the *near* future as it has to do with technology, technologists, and society — between now and, say, 2025 — a little less than ten years from now. In brief, I want to explain why I am not convinced that a technological apocalypse is "nigh at hand."[16] In addition, I cannot refrain from giving a little advice about the future. As the wise Duc de la Rochefoucauld said: "Old people like to give good advice, since they can no longer set bad examples."[17]

Outside the Rose Motel, Easter 1960

At the Big Rock Candy Mountain, about 1961-1962

First, a personal aside. I was the next-to-youngest in my family — my brother Scott and I were called "the little boys." Scott and I were best buddies then, and have been ever since. We were spoiled by our older brother and sisters.

My sisters defended my quirks. My mother engendered in me a love of learning and of the Gospel. Though she never had the chance to attend college, she became very well educated woman through her wide reading on many subjects, including the scriptures and the "best books"[18] of doctrinal and apologetic literature.

My father was the first missionary and the first college graduate in his immediate family. He lovingly spurred me along in my scientific interests, providing me with a large cardboard box in the garage filled with motors, wind-up clocks, and vacuum-tube electronics that I could take apart and *sometimes* put back together.

19

One day at lunch, during the cold war years, after hearing me protest that I didn't want eat my orange because I didn't like the taste of the bitter white stuff that stuck on after you peeled it, my Dad told me that the white stuff would protect me from radiation in the event of nuclear fallout. That creative fib not only solved the immediate problem in getting me to eat my orange but also convinced me that there was a real practical value to science.

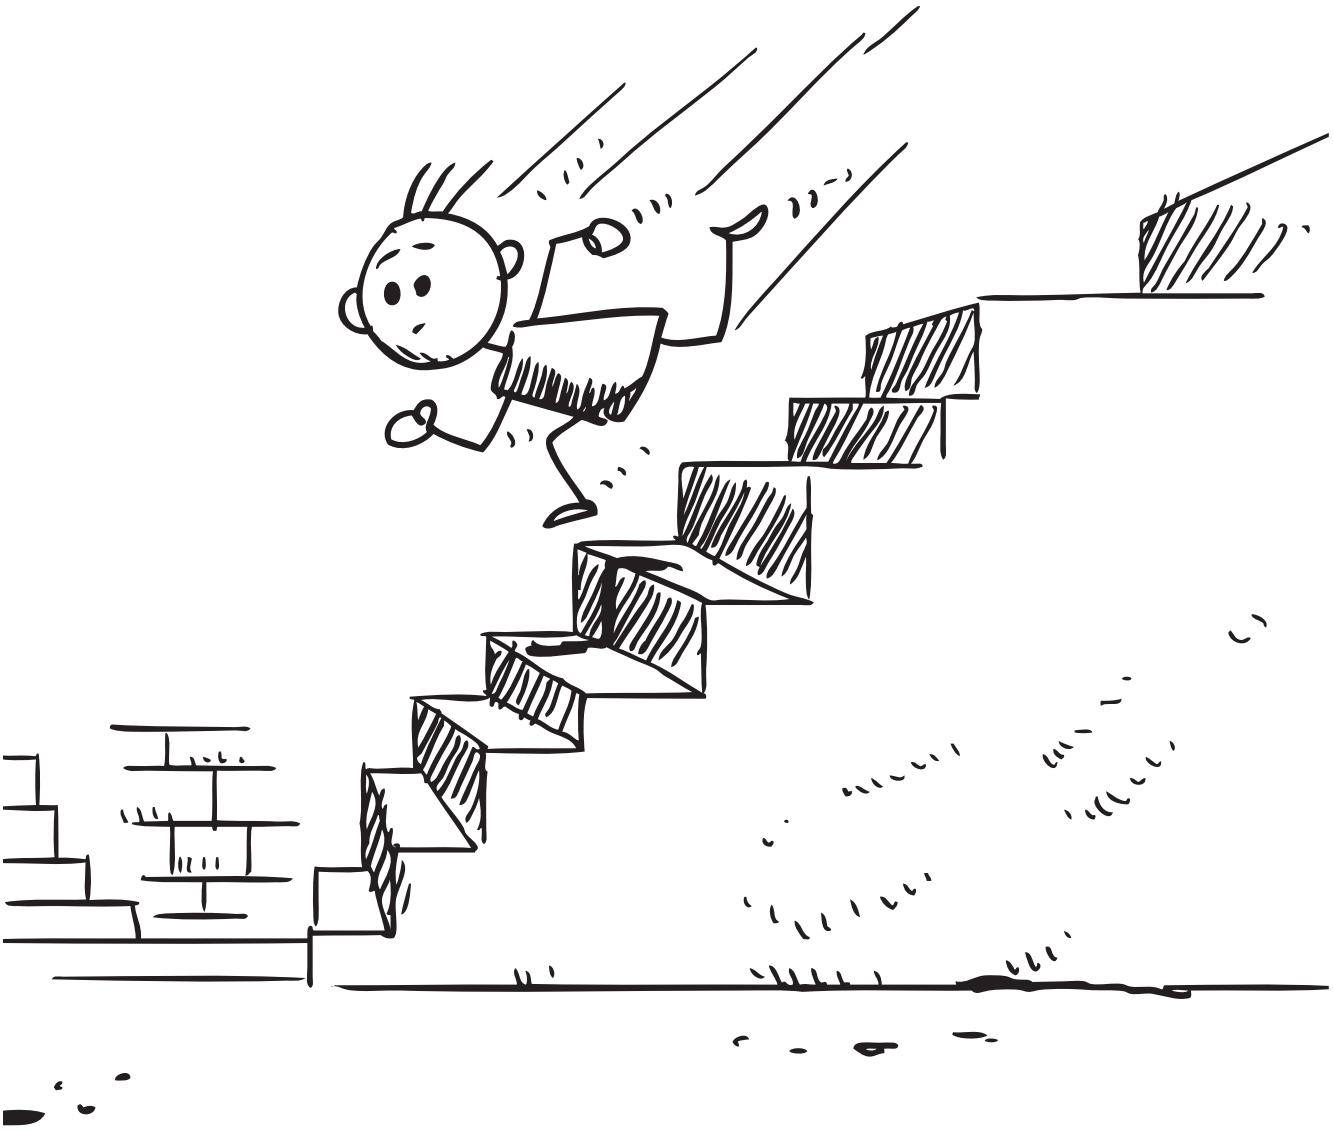

One day my older brother, who I revered then and revere now, convinced me that if I ran fast enough, I could be upstairs and downstairs at the same time. As I recall, my efforts to find out for myself whether that was true lasted long enough to provide amusement for all the family. Though I can't say I succeeded in proving my brother's hypothesis at the time, I was pleased when I later learned enough about quantum mechanics to vindicate my earlier, failed experiments.

I mention all this to convince you that my reflections today come out of a lifetime of watching how scientists and technologists work — sharing from personal experience both their inspiring dreams and their hopeless fantasies.

My day job is at the Florida Institute for Human and Machine Cognition, or IHMC for short.[21] At IHMC, I feel very fortunate to work among researchers who are among the best in the world in their areas of specialty.

One of our most exciting moments in 2015 was the final phase of the DARPA Robotics Challenge.[22] There, the IHMC team walked away with a million-dollar check and top honors among all participating American universities and research institutions, and all but one international competitor. Researchers at IHMC are passionate about science and technology.

However, for the next few minutes, I'd like to share a few of the lessons I've learned, not about science and technology, but rather about *scientists* and *technologists*. If you understand scientists and technologists, you will be able to do a pretty good job in analyzing the news about science and technology, even if you do not understand the science and technology itself.

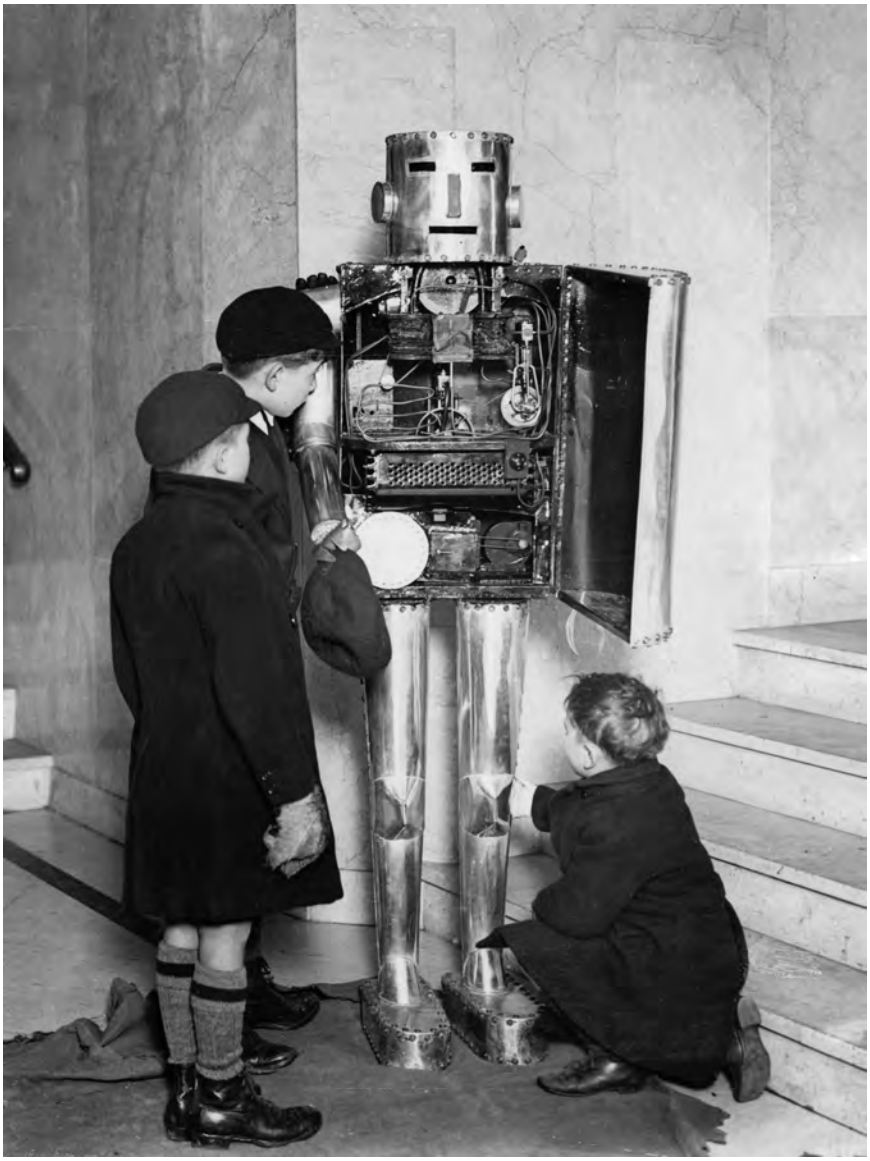

"The original 1930s caption, it must be said, reflects the writer's faith in progress rather than in reality: 'This steel man is near enough to accuracy to explain the physiology of the human frame.'[23]

One thing I have learned about technologists is that they tend to be incurable optimists. For instance, consider the original caption on this photograph from the 1930s: "This steel man is near enough to accuracy to explain the physiology of the human frame."[24] The two students at left are no doubt counting the number of ribs to make sure they are all there. This reminds me of a question I once heard posed about why young Primary children are often asked draw pictures of their fathers to show them on Fathers Day. The answer? Because the fathers want to know what they really look like.

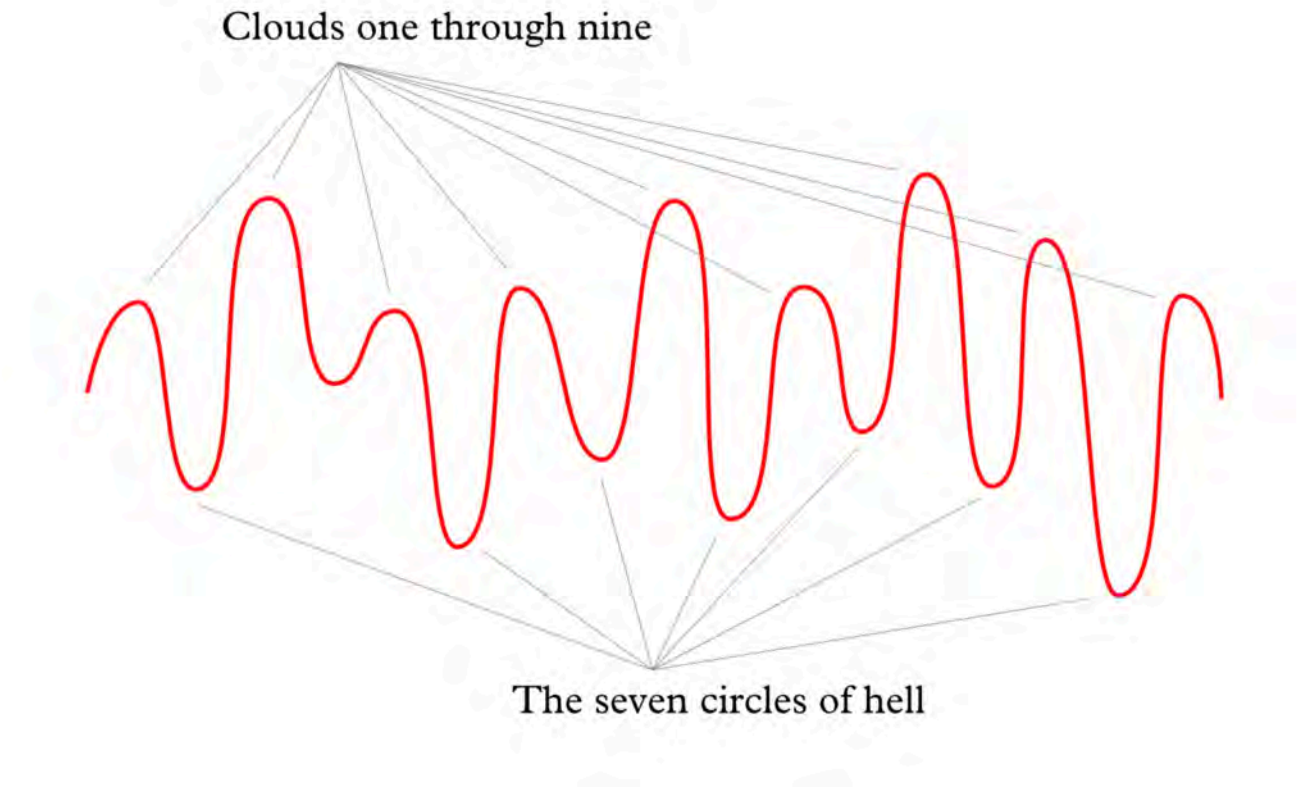

25

When you look at the ups and downs in the lives of scientists and technologists, as shown here, it becomes understandable why they might suffer from incurable optimism, even when the face of reality glares at them fiercely between the eyes.[26] As is often expressed, with unfeigned sympathy, "There goes another beautiful theory about to be murdered by a brutal gang of facts."[27] The professional lives of researchers are inherently unstable, and in many cases their stubborn, unreasonable optimism gives them courage to engage in tedious — and often discouraging — work every day.[28]

The stress of scientists and technologists is not merely a stress of the first order, like the kind that stems from high pressure and overwork, but also of the second order, which is something more existential in nature. Many researchers are passionate about the potential of their contributions, wanting to make a difference in important contemporary problems such as health, poverty, food production, and quality-of-life. Their stress is much like that of the struggling artist or of someone in the early, uncertain stages of a romantic relationship, "who really, really, really want[s] it to work, but lack[s] a clear model of *how*."[29] It is not just a matter of bulldozing one's way to success by working incredibly hard or of becoming more and more sure through experience about "about which path to take, but [rather] about [the uncertainty of] whether the paths (and the destination!) are even real."[30]

And then, if it weren't enough to be grappling with uncertainties relating to the *scientific* viability of the work, researchers usually have to be concerned just as deeply with the dizzyingly frequent changes in levels of public interest in the relevant ideas, which in turn drive the up and down trends of highly competitive funding in their particular lines of investigation.

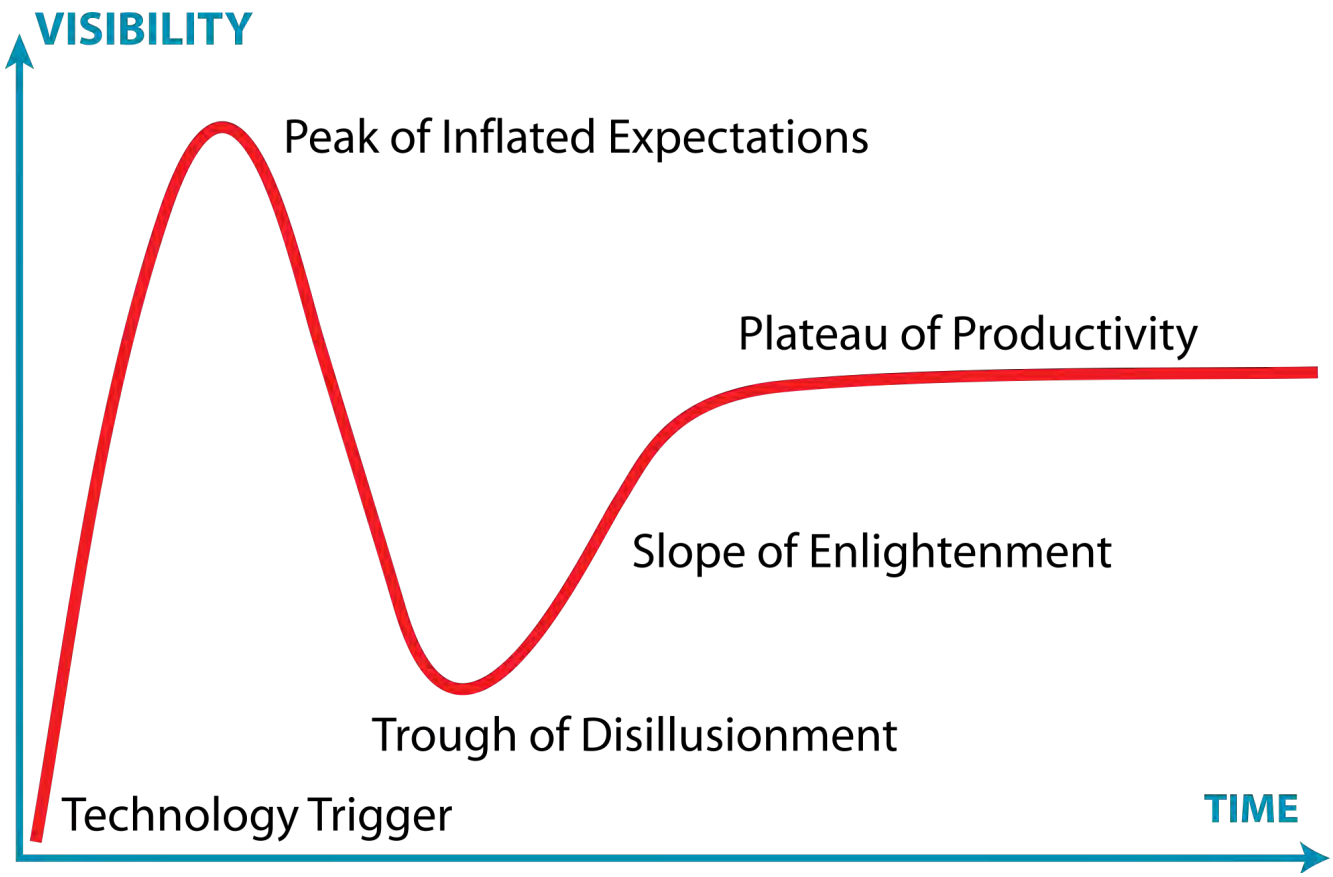

The Gartner Hype Cycle.[31]

Each year the Gartner Group publishes an annual update to their hype cycle.[32] The curve illustrates how the most successful emerging technologies rapidly trend upward toward a peak of "inflated expectations" before suddenly dropping down into a "trough of disillusionment." Only a relative few such technologies sustain their popularity long enough to progress through a gradual "slope of enlightenment" and on to a "plateau of productivity."

Bad timing with respect to the hype cycle can be more destructive to the odds of success in executing a line of research than having a bad idea to begin with. The truth of this claim is evident in the sheer volume of bad ideas that are funded as soon as a given topic approaches the peak of the hype cycle.

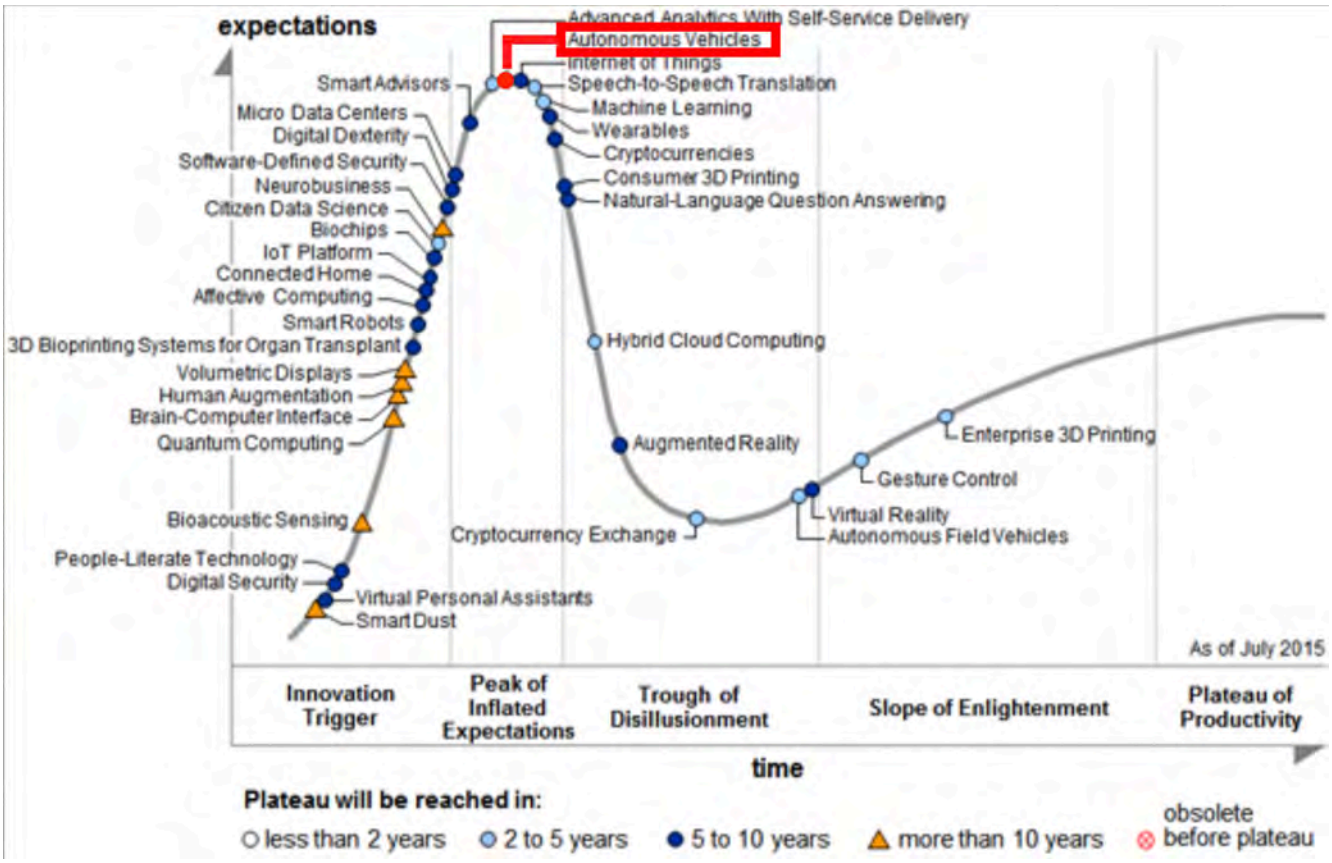

The 2015 Gartner Hype Cycle Update [33]

This brings us to our first example, which sits at the very top of the 2015 Gartner hype cycle update: self-driving cars.

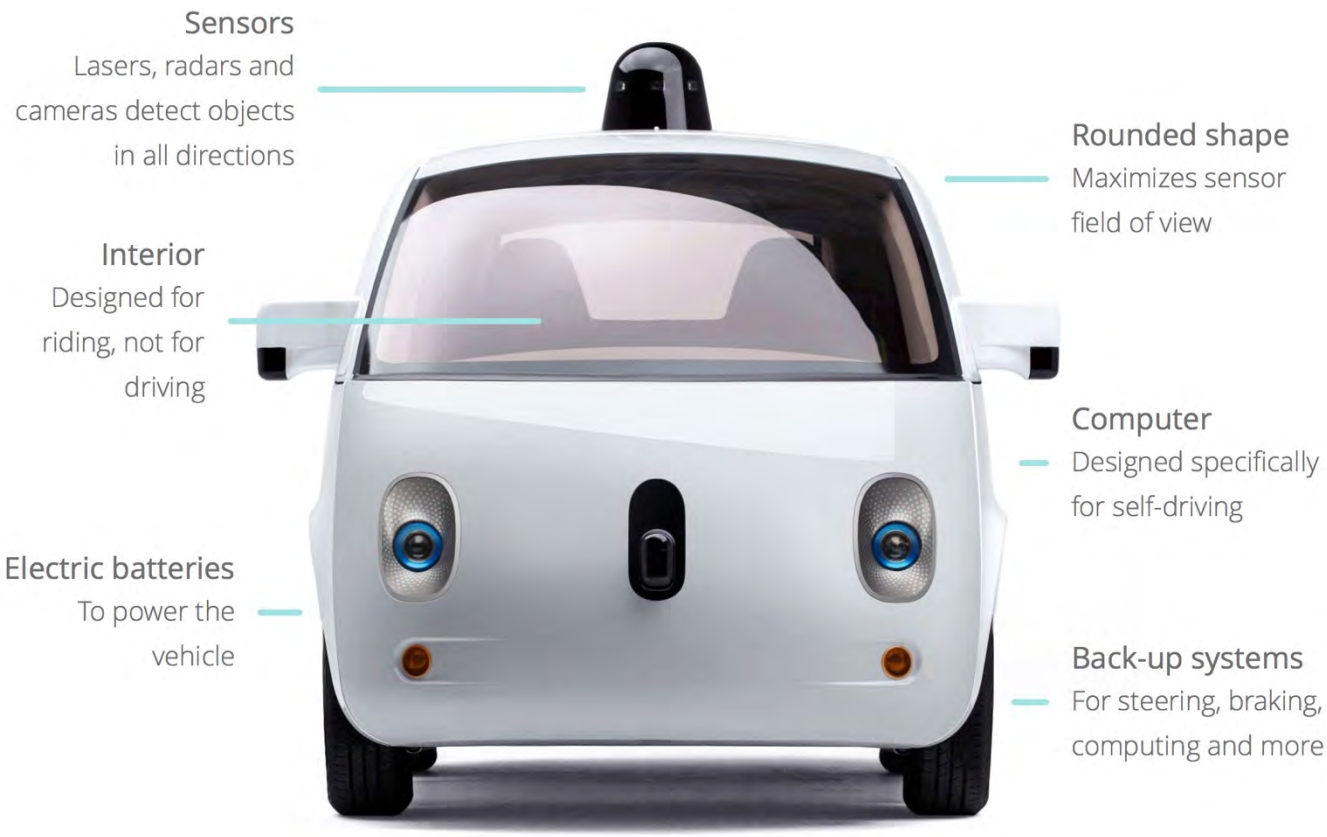

Google Car Prototype[34]

Self-Driving Cars

This is the most easily recognized self-driving car today, the Google Koala prototype that was publicly unveiled in the fall of 2015. I admire the courage of pioneers such as Sebastian Thrun whose tireless research and fearless advocacy of self-driving cars at Google both created groundbreaking technologies and opened up what will surely prove to be one of the biggest transporation developments of the coming century.

The image tells you about the features of the Google car but it does not tell you what features of an *ordinary* car that have been removed, namely, the steering wheel, the gas pedal, and the brake pedal.[35] It is a specific example of the incurable optimism of technologists.

Success in fielding large numbers of *general-purpose* self-driving cars meaning cars that are intended to successfully negotiate the vast majority of situations that manually driven cars do today as opposed to cars that operate in specific, well-constrained niches — depends on solving several difficult problems. The biggest challenges are not in the basics of autonomous driving — getting from A to B. The devil is in the myriad details of unexpected events that can occur while driving.[36]

37

My colleague on the Nissan Science Advisory Council,[38] Donald Norman says it this way: "We know two things about unexpected events: first, they always occur, and second, when they do occur, they are always unexpected."[39]

The California driving authority is trying to do something about this problem. It has recently implemented new rules whereby autonomous vehicles are required "to have means whereby a person sitting in the car could intervene at any time, if the technology fails."[40] Sounds like a good idea, right?

41

Here's the rub: what Norman calls "halfway automation" or what other researchers sometimes call "the handoff problem." In defense of Google's apprehensions, "*halfway* automation" is sometimes a much bigger problem than *full* automation:[42]

> I once argued[, writes Norman,] that the current state of automation was fundamentally unsound because it was in the dangerous middle ground, neither fully automated nor fully manual. Either have no automation or full automation, I argued, but what we have today is halfway automation. Even worse, the system takes over when the going is easy and gives up, usually without any warning, when the going gets tough — just the reverse of what you would want. …
>
> If one cannot automate fully, then the automation that is possible must be applied with great care, sometimes not being invoked, sometimes requiring more human participation than is really needed in order to keep the human drivers informed and attentive.
>
> Full manual control of automobiles is dangerous. Fully automatic control will be safer. The difficulty lies in the transition toward full automation, when only some things will be automated, when different vehicles will have different capabilities, and when even the automation that is installed will be limited in capability. I fear that while the partial automation of driving will lead to fewer accidents, the accidents that do happen will be greater in magnitude, involve more cars, and exact a higher toll. The joint relationship between machines and their humans must be approached with caution.[43]

Will Nissan Beat Google and Uber to Self-Driving Taxis?

By Mark Harris
Posted 26 Feb 2015 | 15:43 GMT

Photo: Nissan

These Nissan taxis are conventional gasoline-powered vehicles, driven by humans. But Nissan is collaborating with NASA to develop a fleet of electric autonomous cabs.

Who will build the first robot taxis? Google has a working prototype but no experience in manufacturing cars (http://spectrum.ieee.org/cars-that-think/transportation/self-driving/googles-selfdriving-car-pals-revealed). Uber, meanwhile, knows the transportation business but has only just started working on autonomous vehicles (http://spectrum.ieee.org/cars-that-think/transportation/self-driving/uber-turns-from-google-teams-up-with-carnegie-mellon-on-selfdriving-cars) with Carnegie Mellon University.

Documents obtained by *IEEE Spectrum* suggest the first cab capable of driving itself (and that you won't feel obliged to tip) might be made by Nissan. In January, the Japanese automaker announced (http://spectrum.ieee.org/cars-that-think/transportation/self-driving/nasa-and-nissan-chase-selfdriving-car-technology) that it would be working with NASA to "demonstrate proof-of-concept remote operation of autonomous vehicles for the transport of . . . goods . . . and people." Using a California Public Records Act request, *Spectrum* has uncovered more details on the particular technologies Nissan and NASA plan to share and, more important, that the main goal of their collaboration appears to be the development of a fleet of remotely-supervised autonomous taxis.

44

The incurable optimism of researchers must be handled with extreme deftness and skill by traditional auto company CEOs. This is not only because they need to temper public expectations but also because many are hoping for partnerships with technology companies such as Google and Apple.[45] In a 2015 article entitled "Will Nissan beat Google and Uber to self-driving taxis?" Nissan's interest in research and development of fleet management services for autonomous vehicles was leaked to the public through a California Public Records Act request.[46]

47

On January 6, 2016, our research team participated with our colleagues at NASA and Nissan in a demonstration of Nissan autonomous driving technologies. The rain made certain aspects of driving difficult, but overall it was a very successful day.

"Carlos Ghosn, chairman and CEO of Nissan Motor, Co., speaks next to a prototype of the autonomous driving Nissan Leaf at Renault-Nissan Silicon Valley in Sunnyvale, California,"[48] January 7, 2016.

The next day, *The New York* Times reported Nissan chairman and CEO, Carlos Ghosn's announcement that Nissan:[49]

> would introduce ten new autonomous vehicles in the next four years.
>
> Elon Musk, the chief executive of Tesla, upped the ante. In a conference call with reporters …, he asserted that the so-called Autopilot feature introduced in the Tesla Model S last fall was "probably better than a person right now."
>
> Mr. Musk also said that within a year or two, it would be technically feasible to summon a Tesla from the opposite side of the country.
>
> But[, continued the *Times*,] there is a growing gap between what these executives are saying and what most people think of when they hear executives or scientists describing … driverless cars.
>
> What Mr. Musk and Mr. Ghosn are describing — cross-country-driving hyperbole aside — are cars with advanced capabilities that can help drive or even take over in tricky situations like parallel parking on a busy street.
>
> Truly autonomous cars that do all the work, like the bubble-shaped vehicles Google has been testing near its Silicon Valley campus, are still at least a decade away from ferrying people around town.

Now I'd like to say a few words about the most incredible example today of the incurable optimism of researchers, namely the building of what has been termed "superintelligence."[50]

Superintelligence

Tremendous progess in our imaginings about superintelligence has taken place in my lifetime. When I was a child, it was too far-fetched to think that anyone could actually *build* a superintelligence, so the best that science fiction could offer us was to help us imagine a *real* human brain, kept alive in a jar and tethered with wires, that was bent on either controlling or destroying the world. Thanks to the broadening of our imaginations in the computer age, we have substituted the outmoded idea of a real brain in a jar with two new and improved substitutes that have become the subject of countless blockbuster films: 1) the omniscient supercomputer, a completely artificial brain;[51] and 2) the omniscient mind, a natural human brain that has been uploaded to a network of supercomputers. Both of these new options for superintelligence — and a few others besides — are being hotly pursued by researchers.

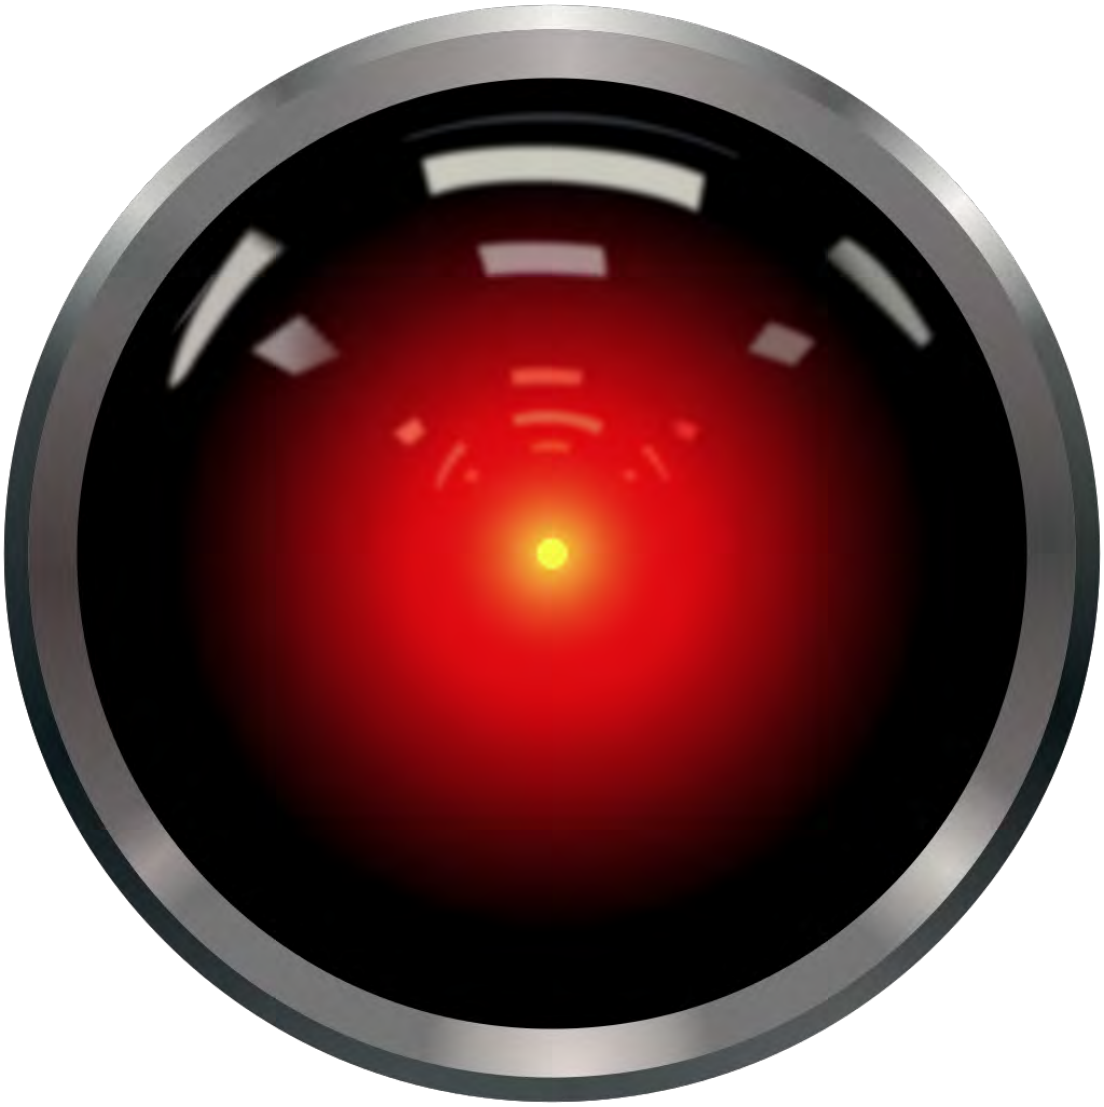

52

As to the omniscient supercomputer, the current front-runner is IBM Watson, which shares conceptual genes with Arthur C. Clarke's HAL 9000 from *2001: A Space Odyssey.* Although transposing the letters H-A-L one letter forward produces I-B-M, any deliberate connection was adamantly denied by Clarke, though later embraced by IBM.[53]

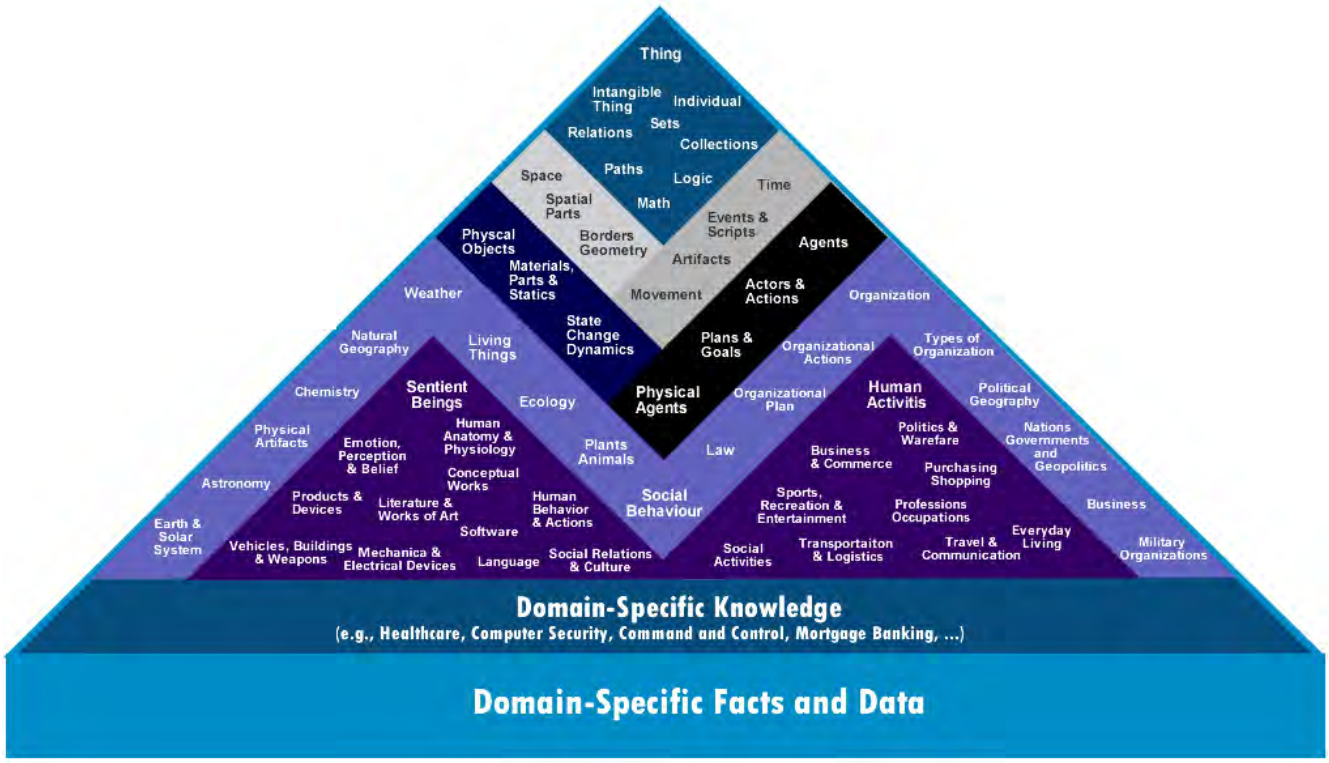

From a research perspective, Watson shares conceptual genes with Doug Lenat's CYC project, an ambitious multi-decade project to build a general purpose AI that has failed to yield the fruits its originators have always dreamed of.[55] However, unlike the current version of CYC, for which bits of knowledge usually have been crafted by hand, Watson has the advantage of being able to ingest large swaths of the Internet and build a knowledge base largely on its own.

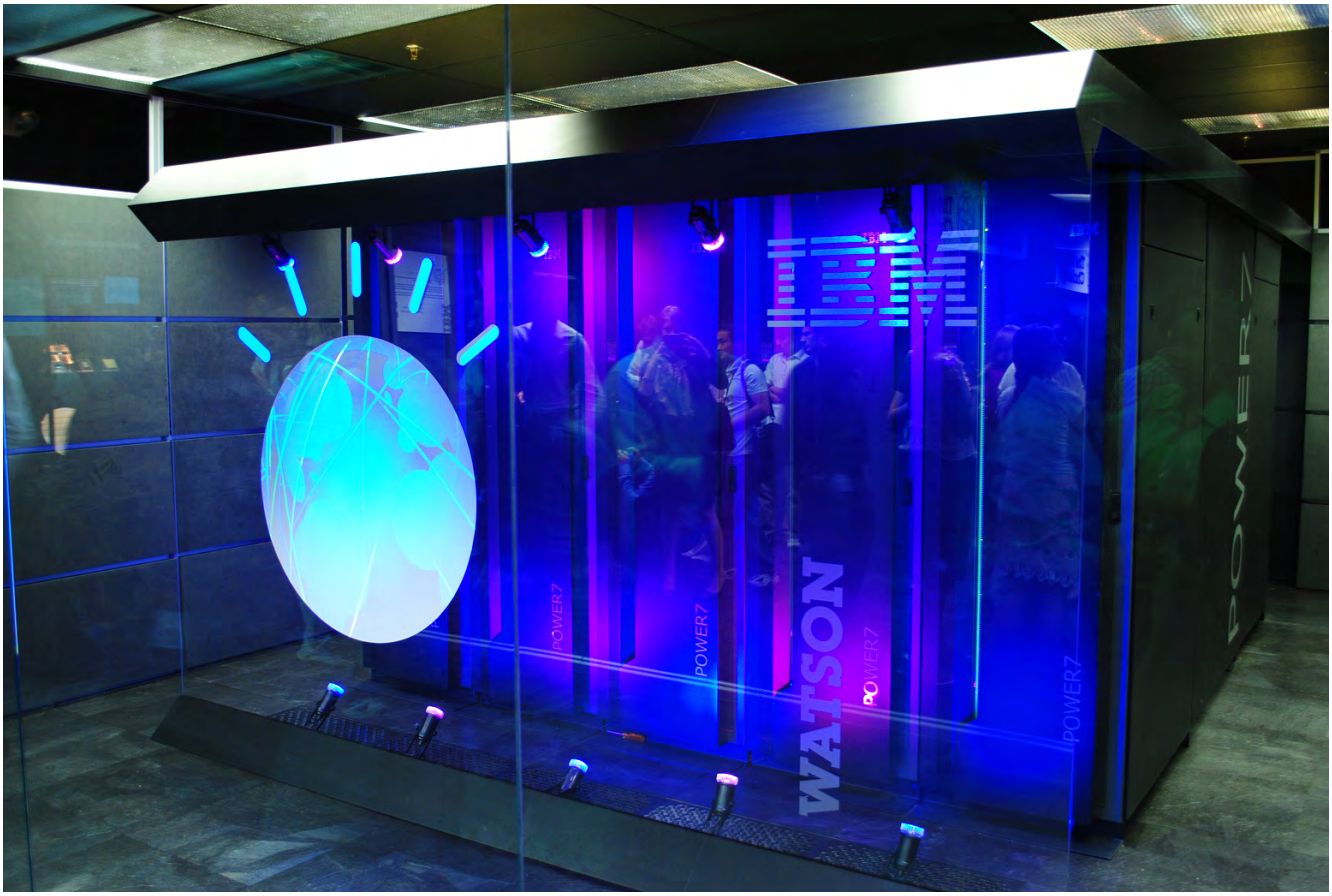

56

To the disappointment of IBM, Watson has not taken off with the speed and glory that they initially hoped since its public debut on *Jeopardy!* in 2011. Indeed, *Jeopardy!* has been the *only* application for Watson that has made much of a splash with the public. IBM's website currently lists only few dozen small company application partners, and a February 2016 article touts with unabashed optimism "future potential" much more than it parades details of its current successes over the last five years.[57] As a super-smart search engine, a capability for complex classification or diagnosis problems, or a natural-language-based analytic assistant, it has great potential. As a superintelligence that matches the ambitions of HAL, I predict it will continue to fall short for the foreseeable future. The top researchers at IBM must already know this; though the sales and marketing folks still seem to be in denial.

58

Another kind of alternative was introduced in a series of books by Ray Kurzweil, the most well known of modern-day transhumanists. He has achieved notoriety for his technology predictions, which he claims have been accurate 86-95% of the time.[59] However, the discrepancies between his self-assessments and the assessment of others of his accuracy resemble the differences between Donald Trump's claims about Trump University and the claims of everyone else.[60] As one critic concludes: "On close examination, [Kurzweil's] clearest and most successful predictions often lack originality or profundity. And most of his predictions come with so many loopholes that they border on the unfalsifiable."[61]

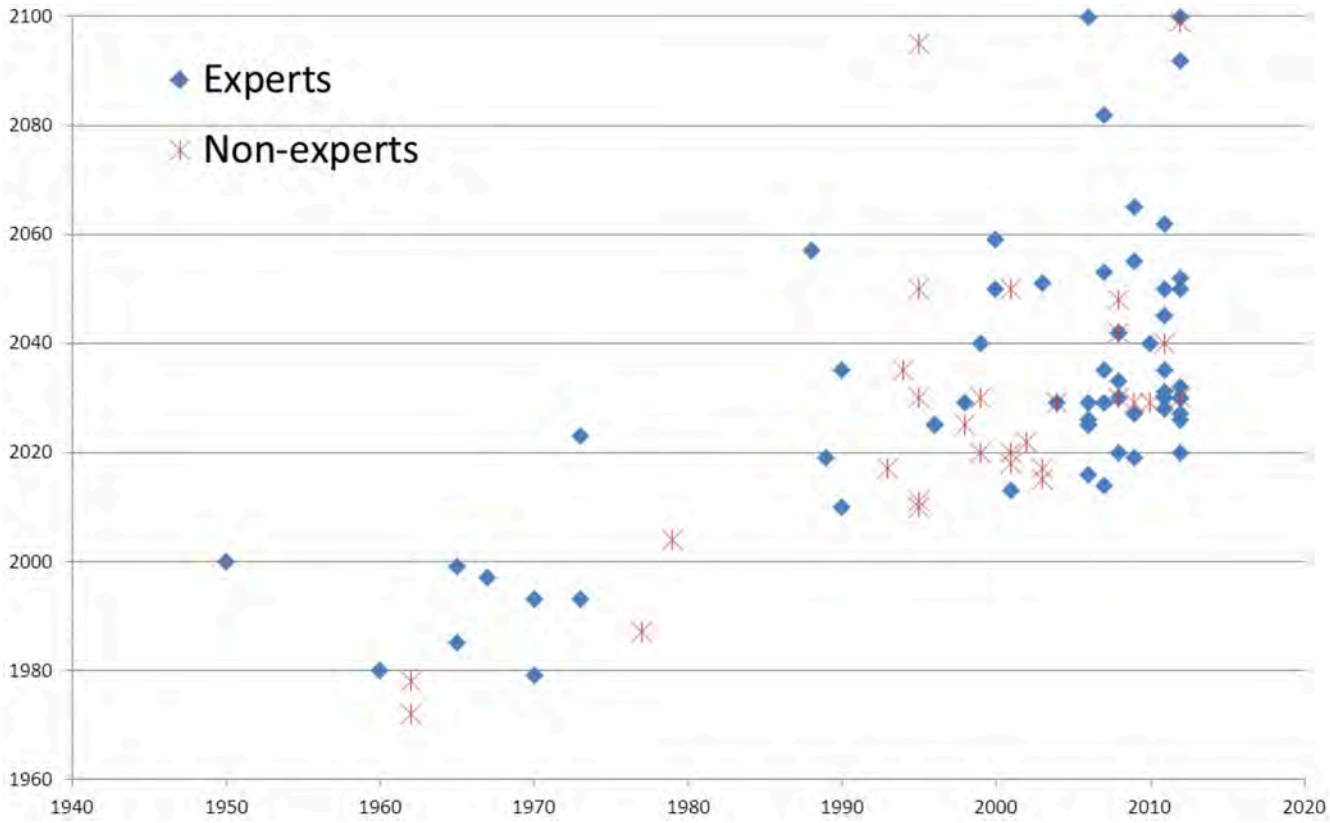

"Median estimate for human –level AI, graphed against date of prediction. Based on data gathered by Wang and Potter for the Singularity Institute, now Machine Intelligence Research Institute (MIRI)."[62]

In Kurzweil's defense, other experts fare equally poorly in their predictions about the future of AI, differing little from the opinions of non-experts in what they say or how accurate they are.[63] For example, this graph shows experts' and non-experts' median-estimates for when "human-level" AI will appear, graphed against the date of prediction.[64] The predictions of both experts and non-experts are all over the map.

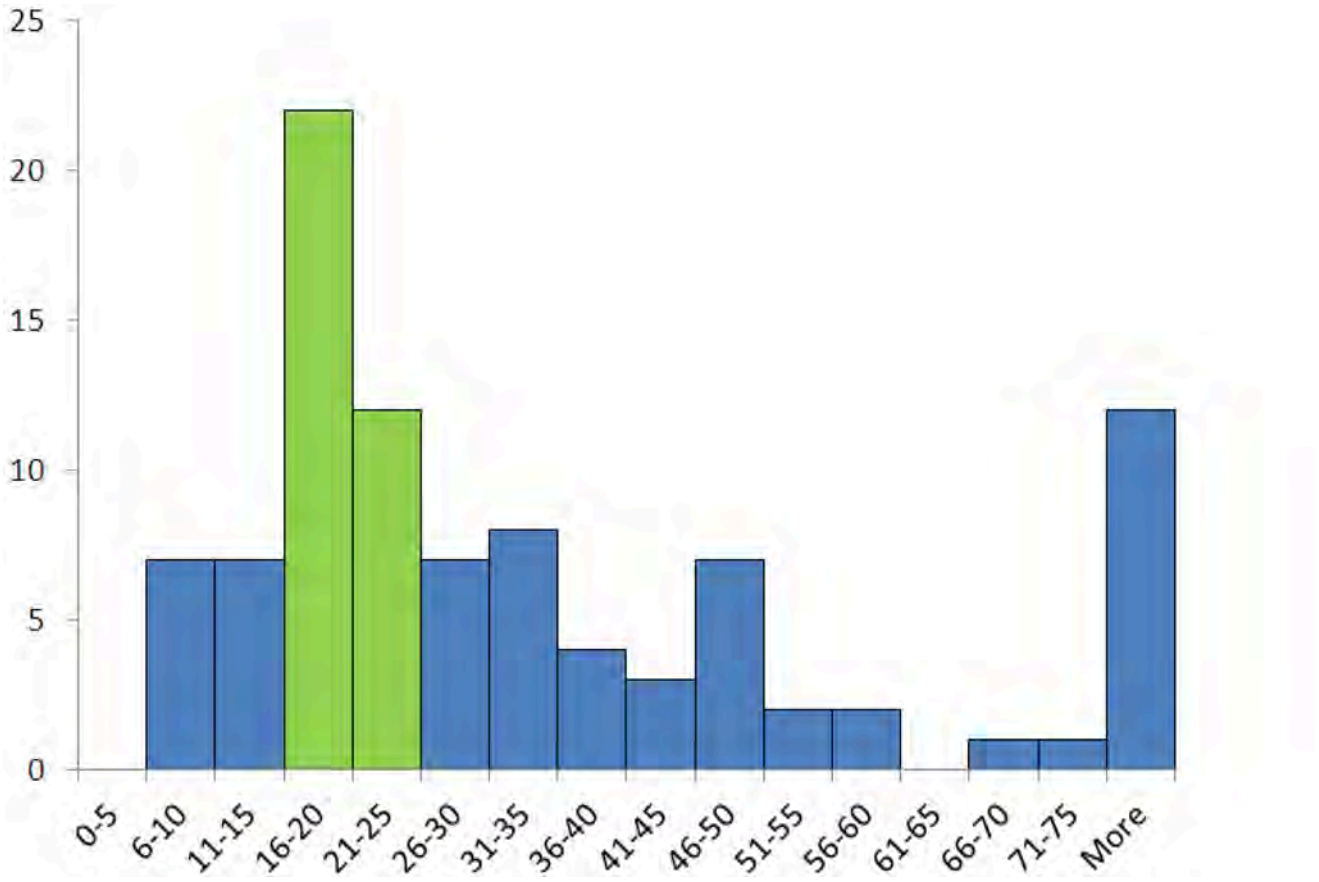

"Time between the arrival of AI and the date the prediction was made. Years on the *x* axis, number of predictions on the *y* axis. Based on data gantered by Wang and Potter for the Singularity Institute, now Machine Intelligence Research Institute (MIRI)."[65]

Notably, "there is a strong tendency to predict AI within 15 to 25 years[, regardless of] when the prediction is made."[66]

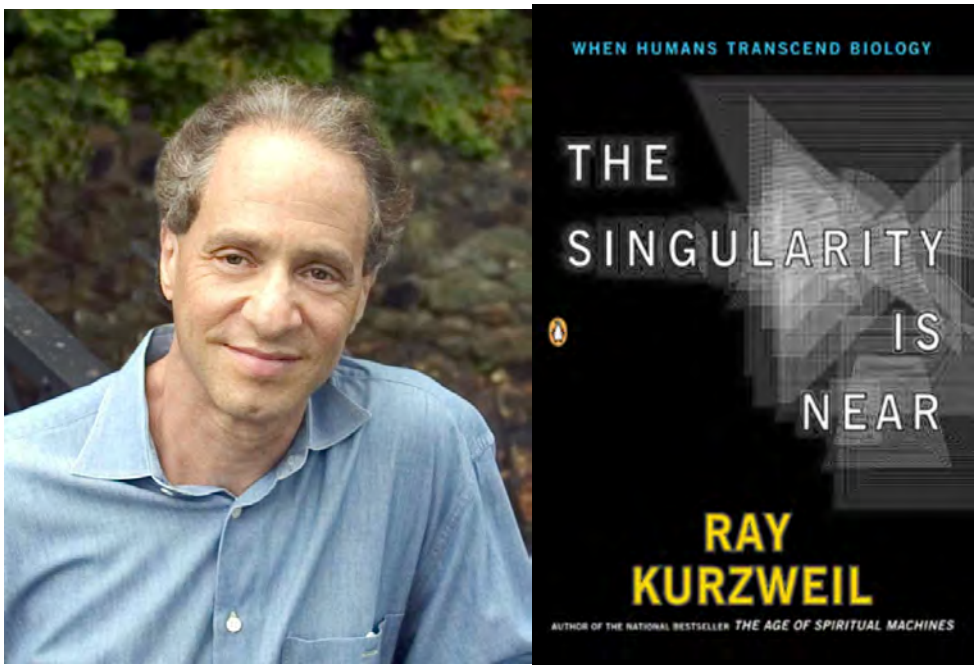

In Kurzweil's 2005 book, *The Singularity is Near: When Humans Transcend Biology*,[67] he predicts a future where, in "[f]using themselves with machines, humans can leave the flesh behind."[68] In the meantime, "set[ting] out a plan of diet, exercise, vitamin supplementation, and preventive medical care" will, he believes, "enhance longevity to the point where technology can overcome mortality."[69] In other words, Kurzweil is doing everything he can to live long enough in mortality so that he can make himself *immortal* through technology. Overall, John Gray sees Kurzweil's program being "best understood as a version of process theology."[70] "It is not essentially different from Gorky's fantasy of humans evolving to become pure thought. … The virtual afterlife is a high-tech variant of the Spiritualist Summerland, while accelerated evolution in cyber-space is an updated version of Myers' Victorian dream of progress in the after-world."[71]

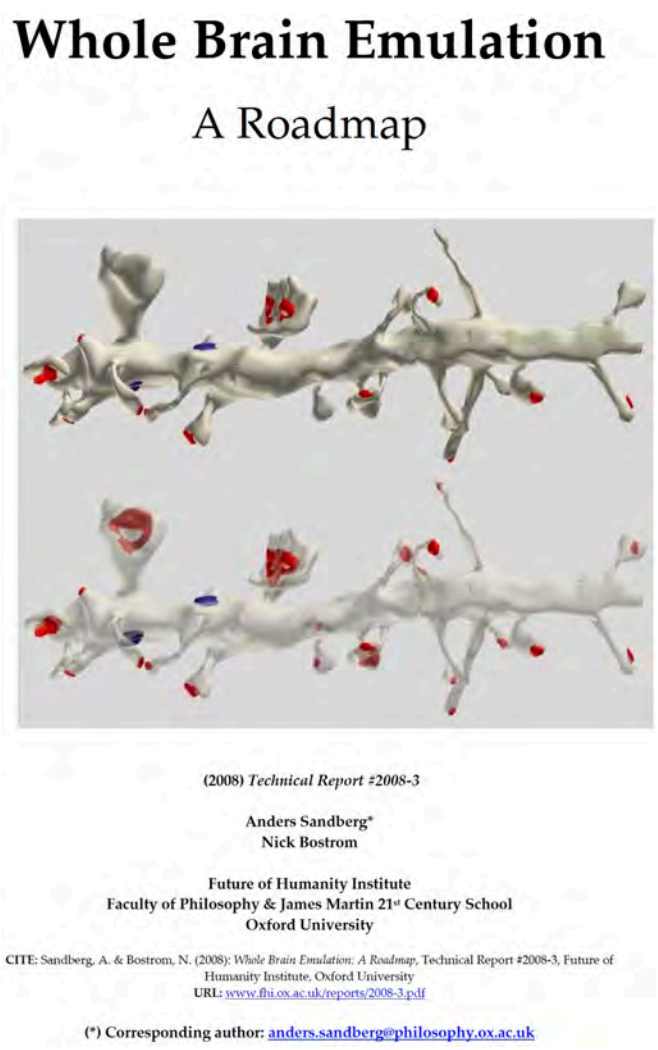

Whole Brain Emulation

A Roadmap

(2008) *Technical Report #2008-3*

Anders Sandberg*
Nick Bostrom

Future of Humanity Institute
Faculty of Philosophy & James Martin 21st Century School
Oxford University

CITE: Sandberg, A. & Bostrom, N. (2008): *Whole Brain Emulation: A Roadmap*, Technical Report #2008-3, Future of Humanity Institute, Oxford University
URL: www.fhi.ox.ac.uk/reports/2008-3.pdf

(*) Corresponding author: anders.sandberg@philosophy.ox.ac.uk

But apart from the philosophical and theological overtones (including, for LDS believers, the question about how the spirit, mind and body relate), what can be said about the scientific feasibility of uploading our physical brains to a computer? There has been some credible thinking on this topic (as well as some very unsound proposals), but nearly every expert on the topic agrees that the many remaining challenges would not be overcome "in the near future."[73] Before you can upload a brain, you need to be able to model it — and most of the focus to date has been on replicating *structure*, ignoring the equally or perhaps more daunting challenge of replicating the complex details of *functions* and *processes*. As part of a 2008 National Academies study of the future of cognitive neuroscience and related technologies, I wrote the following summary of our consensus on the topic:[74]

> [Despite the impressive increases in high-end computing], there does not yet exist either an adequate and detailed understanding of *how* … modeling [of the human brain] can be done, or a complete model of how the brain interacts with complex regulatory and monitoring systems throughout the body. These and other difficulties make it highly unlikely that in the next two decades anyone could build a neurophysiologically plausible model of the whole brain and its array of specialized and general-purpose higher cognitive functions.[75]

The Pentagon is Nervous about Russian and Chinese Killer Robots

Deputy defense secretary: Russia is preparing for all-robot fighting units.DECEMBER 14, 2015
BY PATRICK TUCKER
The Pentagon is rushing to keep up with Russian and Chinese efforts to develop highly autonomous robots — in Russia's case, ones capable of independently carrying out military operations, deputy defense secretary Robert Work told a Center for New American Security national security forum today.

Work quoted the Defense Science Board's **summer study on autonomy and AI**, which said that the human race stands at "an inflection point" in the development of artificial intelligence. Different nations, he noted, are reacting in very different ways.

"We know that China is already investing heavily in robotics and autonomy and the Russian Chief of General Staff [Valery Vasilevich] Gerasimov recently said that the Russian military is preparing to fight on a roboticized battlefield and he said, and I quote, 'In the near future, it is possible that a complete roboticized unit will be created capable of independently conducting military operations.'"

Autonomous Weapons

One area that is fraught with momentous consequences is the prospect of the proliferation of autonomous weapons, one of many topics that I explored in depth as a member of the 2015 Defense Science Board Summer Study on Autonomy.[77] Because the public version of the study has not yet been released, I cannot comment on details nor share the committee's consensus. However, I want to share my personal thoughts and concerns about the future development and deployment of weapons with autonomous capabilities in both the cyber and the physical domains by adversaries who are not constrained by the principles and ethics that are meant to govern US policies in this arena.

Asimov: Three Laws of Robotics

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

As an aside, over the last fifteen years, our research group has spent significant effort on technological solutions to the problem of policy-based governance of intelligent systems, in the spirit of Isaac Asimov's laws of robotics.[78]

KAoS Policy Services

- IHMC's KAoS Policy Services Framework is the most rich, mature, and fully-featured framework of its kind
  - It has been used in dozens of military and space research and development programs
  - The KAoS core ontology was selected by the NSA Digital Policy Management initiative as the basis for its standards effort for all federal government organizations
  - KAoS was recently adopted by a large international industrial company as the basis of an effort to commercialize a new class of complex, networked scientific devices

We call our digital policy services framework KAoS.[79]

80

Unlike nuclear weapons deployed in World War II, the proliferation of autonomous weapons would not be constrained by a given nation's ability to perform sophisticated refinement of rare elements, but rather is being helped along rapidly by the virtually unlimited capacity for just about anyone to share and duplicate the needed software using worldwide computer networks. Unlike nuclear weapons, the development and proliferation of intelligent weaponry cannot be easily monitored or banned. There is no need to solve the long-term AI problem of general intelligence in order to develop early generations of such weapons — only the development of limited-scope autonomous capabilities that are dedicated to specific purposes.[81]

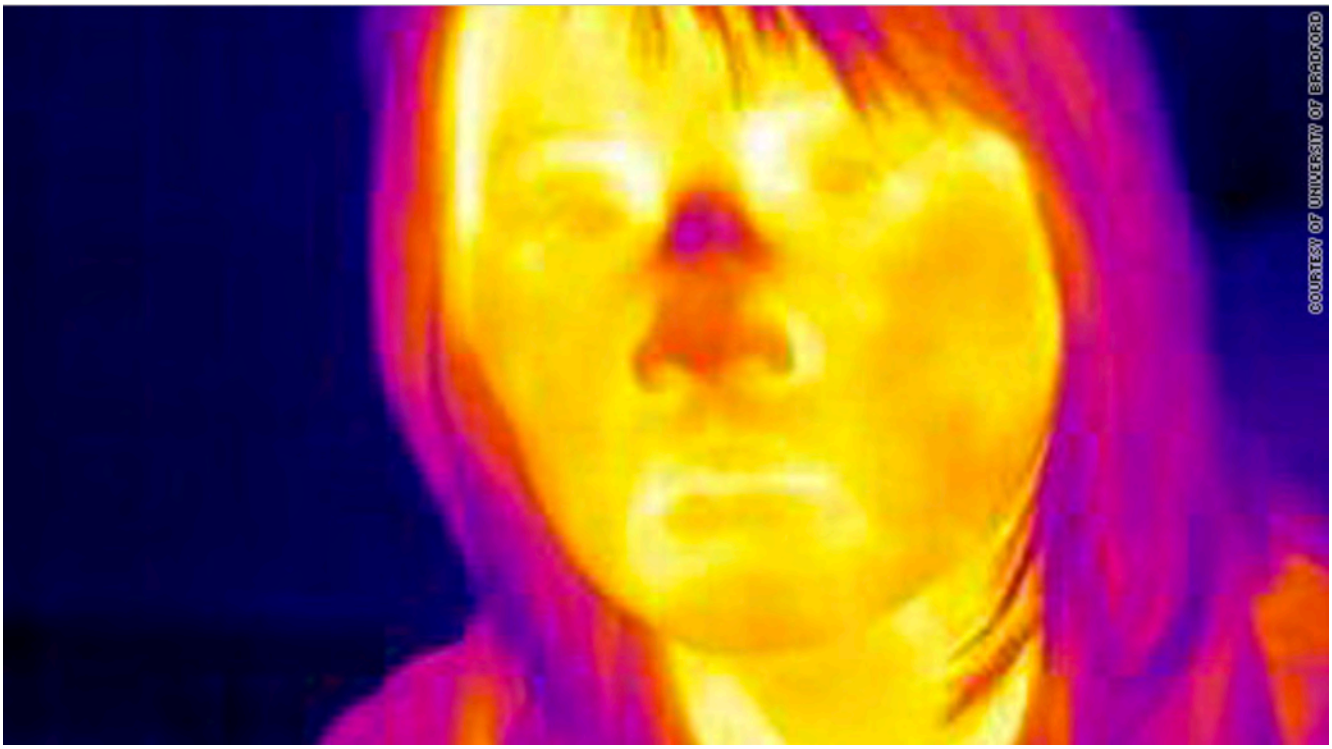

82

Just to prove that others besides researchers can entertain wild speculations, at the initial meeting of a National Academies study a group of us were told that one of the questions we needed to explore was whether it would be possible to develop an autonomous weapon that could fire into a crowd and only hit people with hostile thoughts. Without even entering into the staggering legal and ethical implications of developing such a weapon, our committee implicitly answered this question on the pure grounds of common sense, based on decades of data: Today, we hardly know how to build a good, automatic lie detector, let alone being able to recognize a range of specific psychological states for unknown individuals in an uncontrolled environment — and, thank heavens, it's highly unlikely that the needed breakthroughs will happen anytime in the next few decades.[83] You're safe for now!

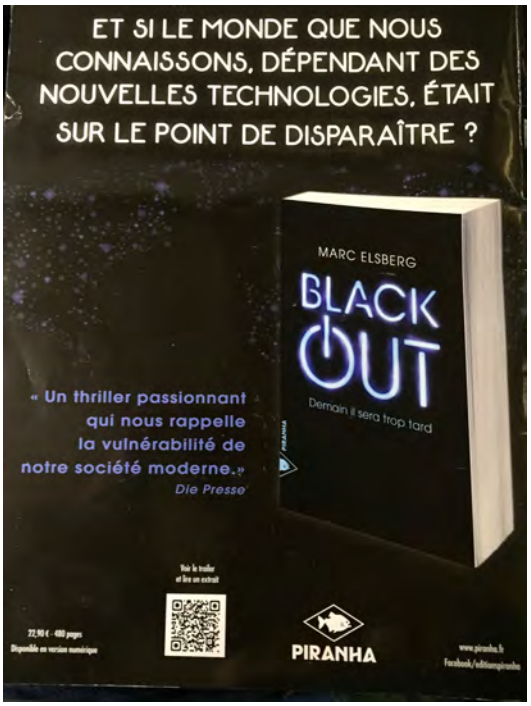

Cyber warfare is one of the most underappreciated threats of the modern age. Everything in our economy, infrastructure, and personal lives would come to a grinding halt were such threats carried out at a large scale.

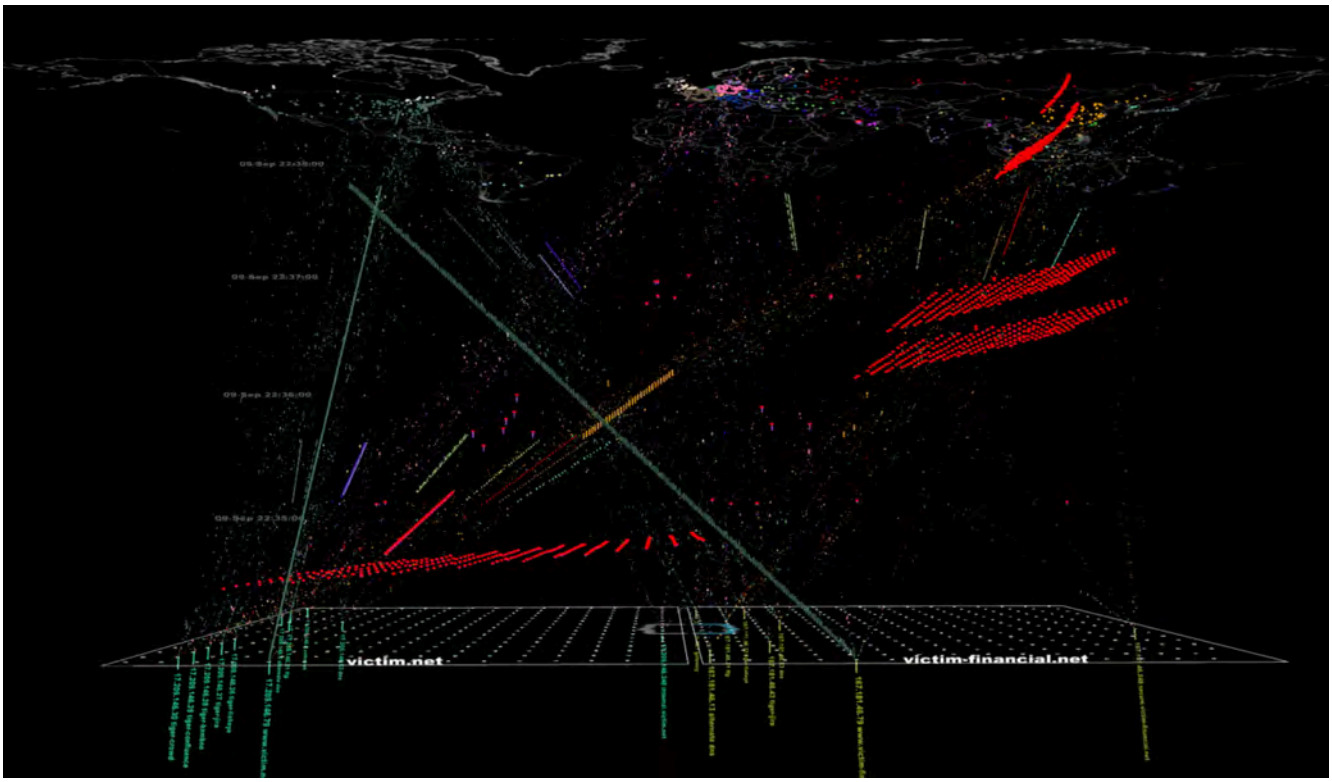

Our Sol cyber framework, here using simulated data, shows one of the approaches our research team developed in response to a government request to be able to visualize and interact with the entire Internet in real time so as to make sense of whatever important events were going on at the moment.[84] We have had a "live," real-time version of such a display continuously working on IHMC's own network for a few years now.

The design of the display exploits specific properties of human perception and cognition, allowing large numbers of interesting events to pop out and be assimilated by the ambient vision system.[85] The visual design is augmented by the capabilities of software agents that work in tandem with analysts to help make sense of the situation.[86] You can see a projection of a world map at the top, with various patterns of attack moving downward toward the company network at the bottom, belonging to a specific victim and its primary financial institution.

While usually rejecting the prospects of an AI explosion, singularity, or apocalypse such as those popularized in the media,[87] researchers in Artificial Intelligence have been thinking more deeply of late about the future of AI. As a result, there has been a recent proliferation of research institutions,[88] studies,[89] articles,[90] books,[91] blogs,[92] and open letters of concern[93] to help assure that both the short- and long-term trajectories of AI research will follow directions that are both safe and beneficial to society. Far from being the neo-Luddites they are sometimes painted to be,[94] these are some of the top minds in the field, believers in the potential of AI for the good of humankind.[95]

Now our brief tour of AI must come to an end. It's been exciting for me over the years to see many of the breakthroughs we used to call Artificial Intelligence become assimilated as ordinary, ho-hum parts of mainstream computer science and engineering.[96] As far as the potential benefits of science and technology, I share much of the optimism of President Gordon B. Hinckley, who said:

> [The twentieth century] has been the best of all centuries. … The fruits of science have been manifest everywhere. … This is an age of greater understanding and knowledge. … This has been an age of enlightenment. The miracles of modern medicine, of travel, of communication are almost beyond belief.[97]

I believe that the fruits of science and technology are divine gifts to which it is appropriate to apply the observation given in D&C 59:20: "And it pleaseth God that he hath given all these things unto man; for unto this end were they made to be used, with judgment, not to excess, neither by extortion."

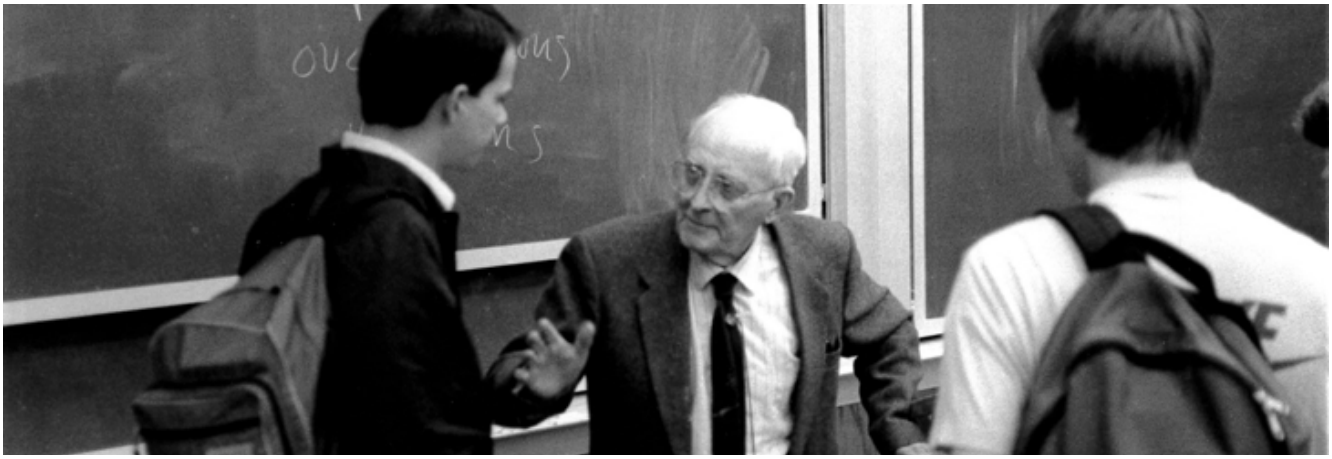

98

Do I ever lose sleep over the *future* of Artificial Intelligence? Only rarely, and that's usually when I'm wrestling with a solution to some interesting problem. That is not to say that I don't sometimes lose sleep over the future in general — for related reasons that are best illustrated by Boyd Petersen's account of an incident involving the late Hugh Nibley:[99]

> One day in the early 1950s, Hugh Nibley's teaching assistant Curtis Wright found Hugh leaning over his desk, reading from the Book of Mormon, and laughing. Wright asked Hugh Nibley what was so funny, and he responded that he had discovered an error in the Book of Mormon. "You did, huh?" Wright asked. "That's interesting. Let me see it."
>
> Hugh handed the scriptures over to Wright and pointed to Alma 42:10, which says that humans are "carnal, sensual, and devilish, by nature." Wright read the passage and demanded, "Well, what's the matter with that?" … Wright was beginning to think that Hugh might be ridiculing the Book of Mormon. "So I got a little defensive," says Wright. Unable to conceal his contempt, Wright demanded, "How's it a mistake?"
>
> He responded, "Well, look at Alma, he says that all mankind is carnal, sensual, and devilish by nature. And he should've said they were carnal, sensual, devilish, and *stupid*."

No, I don't worry too much about the future of Artificial Intelligence, but I do over the consequences of natural stupidity. When Artificial Intelligence meets natural stupidity, unfortunate things can happen. May God grant that we may read and understand the fine print in the hype cycles, discern the "designs which do and will exist in the hearts of conspiring men in the last days,"[100] and, most important of all, rely on divine wisdom to overcome our natural stupidity is my prayer …

References

AAAI Presidential Panel on long-term AI futures: 2008-2009 study. In *Association for the Advancement of Artificial Intelligence*. https://www.aaai.org/Organization/presidential-panel.php. (accessed March 5, 2005).

AI effect. In *Wikipedia*. https://en.wikipedia.org/wiki/AI_effect. (accessed March 12, 2016).

Anonymous. "How IBM makes its partners smarter: IBM's Watson turns cognitive computing into business success." *CyberTrend 14*, no. 2 (February 2016): 10-15. http://www.cybertrend.com/article/19837/how-ibm-makes-its-partners-smarter. (accessed March 5, 2016).

Armstrong, Stuart. 2013. Assessing Kurzweil: The results. In *LessWrong*. http://lesswrong.com/lw/gbi/assessing_kurzweil_the_results/. (accessed March 15, 2016).

Armstrong, Stuart, and Seán ÓhÉigeartaigh. "Who knows anything about anything about AI?" In *Intelligence Unbound: The Future of Uploaded and Machine Minds,* edited by Russell Blackford and Damien Broderick, 46-60. Chichester, West Sussex, England: John Wiley, 2014.

Armstrong, Stuart, Kaj Sotala, and Seán ÓhÉigeartaigh. 2014. The errors, insights and lessons of famous AI predictions — and what they mean for the future. In. http://www.fhi.ox.ac.uk/wp-content/uploads/FAIC.pdf. (accessed March 15, 2016).

Atkinson, Robert D. 2015. The 2015 ITIF Luddite Award nominees: The worst of the year's worst innovation killers. In *Information Technology and Innovation Foundation*. 160315-The Future Isn't What It Used to Be-s.docx. (accessed March 5, 2016).

Autonomous weapons: An open letter from AI and robotics researchers. In *Future of Life Institute*. http://futureoflife.org/open-letter-autonomous-weapons/. (accessed March 5, 2016).

Berra, Yogi. *The Yogi Book: I Really Didn't Say Everything I Said*. New York City, New York: Workman, 1998.

Bostrom, Nick, and Milan M. Cirkovic. *Global Catastrophic Risks*. Oxford, England: Oxford University Press, 2008.

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford, England: Oxford University Press, 2014.

Bradshaw, J. M., Paul Feltovich, and Matthew Johnson. "Human-Agent Interaction." In *Handbook of Human-Machine Interaction,* edited by Guy Boy, 283-302. Ashgate, 2011.

Bradshaw, Jeffrey M., ed. *Software Agents*. Cambridge, MA: The AAAI Press/The MIT Press, 1997.

Bradshaw, Jeffrey M., Marco Carvalho, Larry Bunch, Tom Eskridge, Paul J. Feltovich, Matthew Johnson, and Dan Kidwell. "Sol: An Agent-Based Framework for Cyber Situation Awareness." *Künstliche Intelligenz 26*, no. 2 (2012): 127-40.

Bradshaw, Jeffrey M., and Marco Carvalho. "Multi-agent systems for deep understanding of cyberspace (Abstract of an invited presentation)." Presented at the Eighth Cyber Security and Information Intelligence Research Workshop (CSIIRW 2013), Oak Ridge National Labs, Oak Ridge, TN, January, 2013. http://www.jeffreymbradshaw.net/publications/121102-Multi-Agent Systems Abstract v. 2.pdf. (accessed March 15, 2016).

Bradshaw, Jeffrey M., Robert R. Hoffman, Matthew Johnson, and David D. Woods. "The seven deadly myths of 'autonomous systems'." *IEEE Intelligent Systems 28*, no. 3 (May/June 2013): 54-61. http://jeffreymbradshaw.net/publications/IS-28-03-HCC_1.pdf. (accessed March 15, 2016).

Bradshaw, Jeffrey M., Andrzej Uszok, Maggie Breedy, Larry Bunch, Thomas C. Eskridge, Paul J. Feltovich, Matthew Johnson, James Lott, and Michael Vignati. "The KAoS Policy Services Framework." Presented at the Eighth Cyber Security and Information Intelligence Research Workshop (CSIIRW 2013), Oak Ridge National Labs, Oak Ridge, TN, January, 2013. http://www.jeffreymbradshaw.net/publications/CSIIRW KAoS paper-s.pdf. (accessed March 15, 2016).

Bradshaw, Jeffrey M., Andrzej Uszok, and Rebecca Montanari. "Policy-Based Governance of Complex Distributed Systems: What Past Trends Can Teach Us about Future Requirements." In *Engineering Adaptive and Resilient Computing Systems,* edited by Niranjan Suri and G. Cabri, 259-84. Boca Raton, FL: CRC Press/Taylor and Francis, 2014. http://www.jeffreymbradshaw.net/publications/130529-PolicyDraft2May-jb-rm-jb.pdf.

Bradshaw, Jeffrey M. 2016. Human-Agent-Robot Teamwork Through Coactive Design. Academic Ceremony on the Occasion of the Bestowal of an Honorary Doctorate for Sebastian Thrun. Video Recording of invited Presentation at the TU Delft 174th Dies Natalis Seminar on Robots: Tools or Teammates? (8 January 2016). In *Delft Technical University (TU Delft)*. http://jeffreymbradshaw.net/videos/160108-Bradshaw-Human-Agent-Robot Teamwork Through Coactive Design-TU Delft-174th Dies Natalis Seminar-Robots-Tools or Teammates-edited.mp4. (accessed March 15, 2016).

Bradshaw, Jeffrey M., and Matthew Johnson. 2016. Lessons learned in the theory and practical application of Human-Agent-Robot Teamwork. . Invited presentation at the University of Kaiserslautern, Kaiserslautern, Germany, 17 February 2016. In *University of Kaiserslautern*. http://jeffreymbradshaw.net/videos/Clip2.mov.

Bradshaw to chair new Nissan advisory board. In *IHMC News*. http://www.ihmc.us/bradshaw-to-chair-new-nissan-advisory-board/. (accessed March 15, 2016).

Bunch, Larry, Jeffrey M. Bradshaw, Marco Carvalho, Tom Eskridge, Paul J. Feltovich, James Lott, and Andrzej Uszok. "Human-Agent Teamwork in Cyber Operations: Supporting Co-Evolution of Tasks and Artifacts with Luna." Presented at the Tenth German Conference on Multiagent System Technologies (MATES 2012) (LNAI 7598), Trier, Germany, October 10-12, 2012, 53-67.

Bunch, Larry, Jeffrey M. Bradshaw, Robert R. Hoffman, and Matthew Johnson. "Principles for human-centered interaction design, part 2: Can humans and machines think together?" *IEEE Intelligent Systems 30*, no. 3 (May/June 2015): 68-75. http://www.jeffreymbradshaw.net/publications/30-03-HCC.PDF. (accessed March 15, 2016).

CACM Oct. 2015 - Rise of Concerns about AI: Reflections and Directions. In *Association for Computing Machinery (ACM)*. https://www.youtube.com/watch?v=rRqusoGqyDg. (accessed March 5, 2016).

Clarke, Arthur C. "The future isn't what it used to be." *Engineering and Science 33*, no. 7 (1970): 4-9. http://resolver.caltech.edu/CaltechES:33.7.clarke. (accessed September 18, 2015).

———. "Foreword: The birth of HAL." In *HAL's Legacy: 2001's Computer as Dream and Reality,* edited by David G. Stork, xi-xvi. Cambridge, MA: The MIT Press, 1997.

Clarke, Roger. "Asimov's laws of robotics: Implications for information technology, Parts 1 and 2." *IEEE Computer*, December/January 1993-1994, 53-61/57-66.

Cole-Turner, Ronald, ed. *Transhumanism and Transcendence: Christian Hope in an Age of Technological Enhancement*. Washington, DC: Georgetown University Press, 2011.

Come hear Jeff Bradshaw present on: Agents, ontologies, policies, and visualization: Building blockes for the human-machine interfaces of the future (Presentation, 17 March 2016). In

Silicon Valley Autonomous Vehicle Enthusiasts. http://www.meetup.com/Autonomous/events/229486742/. (accessed March 15, 2016).
Cordeschi, Roberto. "The discovery of the artificial. Some protocybernetic developments 1930-1940." *AI and Society 5* (1991): 218-38. http://philpapers.org/archive/CORTDO-9. (accessed March 12, 2016).
Cyc. In *Wikipedia*. https://en.wikipedia.org/wiki/Cyc. (accessed March 15, 2016).
Dietterich, Thomas G., and Eric J. Horvitz. "Viewpoint: Rise of concerns about AI: Reflections and directions." *Communications of the ACM 58*, no. 10 (October 2015): 38-40. http://cacm.acm.org/magazines/2015/10/192386-rise-of-concerns-about-ai/fulltext. (accessed March 5, 016).
Doyle, Charles Clay, Wolfgang Mieder, and Fred R. Shapiro. *The Dictionary of Modern Proverbs*. Princeton, NJ: Yale University Press, 2012. https://books.google.com/books?id=c0ZiqVnQju8C. (accessed March 15, 2016).
Eliason, Eric A. *The J. Golden Kimball Stories*. Urbana, IL: University of Illinois Press, 2007.
Fife, Austin, and Alta Fife. *Saints of Sage and Saddle*. Bloomington, IN: Indiana University Press, 1956.
The future is not what it used to be. In *Quote Investigator*. http://quoteinvestigator.com/2012/12/06/future-not-used/. (accessed September 18, 2015).
The future isn't what it used to be (Conversations about the future with Arthur C. Clarke, R. Buckminster Fuller, and Alvin Toffler. Directed and produced by Ralph Curtis; writer, Stephan Chodorov. 13 September 1970). In *Camera Three*. http://trove.nla.gov.au/work/22719269?q&sort=holdings+desc&_=1442630718653&versionId=181101751 http://lccn.loc.gov/96521231. (accessed September 18, 2015).
Gage, Fred H., ed. *What is the Brain Good For? Special Issue of Daedalus: Journal of the American Academy of Arts and Sciences 144:1*, 2015.
Gartner hype cycle. In *Wikipedia*. https://commons.wikimedia.org/wiki/File:Gartner_Hype_Cycle.svg. (accessed January 22, 2016.
Gartner's 2015 hype cycle for emerging technologies identifies the computing innovations that organizations should monitor. In *Gartner Newsroom*. http://www.gartner.com/newsroom/id/3114217. (accessed January 22, 2016).
Gray, John. *The Immortalization Commission: Science and the Strange Quest to Cheat Death*. New York City, NY: Farrar, Straus and Giroux, 2011.
Green, Christopher C., Diane E. Griffin, James J. Blascovich, Jeffrey M. Bradshaw, Scott C. Bunch, John Gannon, Michael Gazzaniga, Elizabeth Loftus, Gregory J. Moore, Jonathan Moreno, John R. Rasure, Mark (Danny) Rintoul, Nathan D. Schwade, Ronald L. Smith, Karen S. Walch, and Alice M. Young. *Emerging Cognitive Neuroscience and Related Technologies. A Report of the Committee on Military and Intelligence Methodology for Emergent Neurophysiological and Cognitive/Neural Science Research in the Next Two Decades*. Washington, DC: The National Academies Press, 2008. http://www.nap.edu/catalog/12177/emerging-cognitive-neuroscience-and-related-technologies. (accessed March 5, 2016).
Greenfield, Suan. *Mind Change: How Digital Technologies Are Leaving Their Marks on Our Brains*. New York City, NY: Random House, 2015.
Guha, R.V., and Douglas B. Lenat. "Cyc: A midterm report." *AI Magazine*, Fall 1990, 33-58.

Harris, Mark. 2015. Will Nissan beat Google and Uber to self-driving taxis? In *IEEE Spectrum*. http://spectrum.ieee.org/cars-that-think/transportation/self-driving/will-nissan-beat-google-and-uber-to-self-driving-taxis. (accessed February 27, 2015).

Heisler, Yoni. In 1983 speech, Steve Jobs alluded to the iPad, Siri, the App Store, mainstream Internet connectivity, Google Maps and more. In *Engadget*. http://www.engadget.com/2013/11/14/in-rare-1983-speech-steve-jobs-alluded-to-the-ipad-siri-the-a/. (accessed September 18, 2015).

Hinckley, Gordon B. "Thanks to the Lord for His blessings." *Ensign 29*, May, 88-89.

Hofstadter, Douglas R. *Gödel, Escher, Bach: An Eternal Golden Braid*. New York City, NY: Vintage Books, 1979.

How it works. In *Google Self-Driving Car Project*. https://www.google.com/selfdrivingcar/how/. (accessed March 3, 2016).

Huxley, Thomas Henry, Michael Foster, and E. Ray Lankester. *The Scientific Memoirs of Thomas Henry Huxley*. 4 vols. Vol. 3. New York City, NY: Macmillan, 1901. https://books.google.com/books/about/The_scientific_memoirs_of_Thomas_Henry_H.html?id=9-oEAAAAYAAJ. (accessed February 28, 2016).

Huxley, Thomas Henry. 1870. "Biogenesis and abiogenesis (The presidential address to the British Association for the Advancement of Science for 1870)." In *Collected Essays*. Vol. 8, 229-71. Cambridge, England: Cambridge University Press, 2011. http://ebooks.cambridge.org/chapter.jsf?bid=CBO9781139149273&cid=CBO9781139149273A013. (accessed February 28, 2016).

Jobs, Steve. 1983. The future isn't what it used to be (Presentation to the International Design Conference in Aspen (IDCA), 15 June 1983). In *SoundCloud*. https://w.soundcloud.com/player/?url=http%3A%2F%2Fapi.soundcloud.com%2Ftracks%2F62010118&show_artwork=true. (accessed September 18, 2015).

Johnson, Matthew, Jeffrey M. Bradshaw, Robert R. Hoffman, Paul J. Feltovich, and David D. Woods. "Seven cardinal virtues for human-machine teamwork: Examples from the DARPA Robotic Challenge." *IEEE Intelligent Systems 29*, no. 6 (November/December 2013): 74-80. http://www.jeffreymbradshaw.net/publications/56. Human-Robot Teamwork_IEEE IS-2014.pdf. (accessed March 15, 2016).

Johnson, Matthew, Jeffrey M. Bradshaw, Paul J. Feltovich, Catholijn M. Jonker, M. Birna van Riemsdijk, and Maarten Sierhuis. "Coactive design: Designing support for interdependence in joint activity." *Journal of Human-Robot Interaction 3*, no. 1 (2014): 43-69. http://jeffreymbradshaw.net/publications/173-1278-1-PB.pdf. (accessed March 15, 2016).

Jones, Randolph M., Ryan O'Grady, Denise Nicholson, Robert R. Hoffman, Larry Bunch, Jeffrey M. Bradshaw, and Ami Bolton. "Modeling and integrating cognitive agents within the emerging cyber domain." Presented at the Proceedings of the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC), 20 November-4 December 2015, Orlando, FL 2015. http://www.jeffreymbradshaw.net/publications/15232-iitsec-v5.pdf. (accessed March 15, 2016).

Joseph, Susan, Andrea Westerinen, Jeffrey M. Bradshaw, Peter Sell, and Sandi Roddy. 2012. Digital Policy Management: Be part of the solution, not the problem. In *RSA Conference (28 February 2012)*. http://www.rsaconference.com/events/us12/agenda/sessions/721/digital-policy-management-be-part-of-the-solution. (accessed March 15, 2016).

Kaplan, Mordecai M. "Random thoughts." *The Reconstructionist 16*, no. 2 (March 1950): 30.

Kendall, Frank. 2014. Terms of Reference — Defense Science Board 2015 Summer Study on Autonomy. In *Defense Science Board*. http://www.acq.osd.mil/dsb/tors/TOR-2014-11-17-Summer_Study_2015_on_Autonomy.pdf. (accessed March 15, 2016).

Koene, Randal A. "Feasible mind uploading." In *Intelligence Unbound: The Future of Uploaded and Machine Minds,* edited by Russell Blackford and Damien Broderick, 90-101. Chichester, West Sussex, England: John Wiley, 2014.

Kurzweil, Ray. *The Singularity is Near: When Humans Transcend Biology*. New York City, NY: Penguin Books, 2005.

Lenat, Douglas B., M. Prakash, and M. Shepard. "CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks." *AI Magazine 6*, no. 4 (1986).

Lenat, Douglas B., and R.V. Guha. *Building Large Knowledge-based Systems*. Reading, MA: Addison-Wesley, 1990.

Lenat, Douglas B. "From 2001 to *2001*: Common sense and the mind of HAL." In *HAL's Legacy: 2001's Computer as Dream and Reality,* edited by David G. Stork, 193-209. Cambridge, MA: The MIT Press, 1997.

———. 2015. Computers with common sense. TEDxYouth@Austin (video published 21 May 2015). In *YouTube*. https://www.youtube.com/watch?v=2w_ekB08ohU. (accessed March 15, 2016).

Lindner, Roland. Google warning about autonomous cars. In *Frankfurter Allgemeine (13 January 2016, translated from German)*. http://www.faz.net/aktuell/wirtschaft/netzwirtschaft/google/testfahrten-google-warnt-vor-eigenen-roboterautos-14012386.html (accessed March 4, 2016).

Markoff, John. 2016. For now, self-driving cars still need humans. In *The New York Times*. http://nyti.ms/1ncdLfl. (accessed January 20, 2016).

Maxwell, Neal A. "The stern but sweet seventh commandment (Devotional, Salt Lake Institute of Religion, 8 December 1978)." Salt Lake City, UT: Salt Lake Institute of Religion, The Church of Jesus Christ of Latter-day Saints, 1978. Reprint, Abridged Versions: The stern but sweet seventh commandment, *New Era*, June 1979 (https://www.lds.org/new-era/1979/06/the-stern-but-sweet-seventh-commandment?lang=eng), *New Era*, March 2003 (https://www.lds.org/new-era/2003/03/reasons-to-stay-pure?lang=eng).

McDermott, Drew. "Artificial intelligence meets natural stupidity." *SIGART Newsletter*, no. 57 (1976): 4-9.

Minsky, Marvin, ed. *Semantic Information Processing*. Boston, MA: The MIT Press, 1968.

———. 1961. "Steps toward artificial intelligence." In *Computers and Thought,* edited by Edward A. Feigenbaum and Julian Feldman, 406-50. New York City, NY: McGraw-Hill, 1963.

Moody, Sidney. 1999. The brain behind Cyc: Scientist Doug Lenat discusses Artificial Intelligence. In *The Austin Chronicle*. http://www.austinchronicle.com/screens/1999-12-24/75252/. (accessed March 15, 2016).

Neubauer, Hendrik. *Curious Moments: Archive of the Century (Archives of Text och Bilder, Stokholm, made available through SVT Bild/Das Fotoarchiv)*. Köln, Germany: Könemann, 1999.

New Scientific Advisory Council to assist at Nissan Research Center - Silicon Valley (12 March 2014). In *Nissan News Release*. http://nissannews.com/en-US/nissan/usa/releases/new-scientific-advisory-council-to-assist-at-nissan-research-center-silicon-valley. (accessed March 15, 2016).

New tools counter cyber threats. In *IHMC Newsletter 10:1 (Feburary 2013)*. http://www.jeffreymbradshaw.net/publications/IHMCNewslettervol10iss1.pdf. (accessed March 15, 2016).
Norman, Donald A. *The Design of Future Things*. New York City, NY: Basic Books, 2007.
One hundred year study on Artificial Intelligence (AI100). In *Stanford University*. https://ai100.stanford.edu. (accessed March 5, 2016).
An open letter: Research priorities for robust and beneficial Artificial Intelligence. In. http://futureoflife.org/ai-open-letter/. (accessed March 5, 2016).
Paul Valéry (1871-1945). In *The Poetry Foundation - Poems and Poets*. http://www.poetryfoundation.org/bio/paul-valery. (accessed January 22, 2016).
Petersen, Boyd Jay. *Hugh Nibley: A Consecrated Life*. Draper, UT: Greg Kofford Books, 2002.
Pynadath, David, and Milind Tambe. "Revisiting Asimov's first law: A response to the call to arms." Presented at the Proceedings of ATAL 01 2001.
Ray Kurzweil. In *Wikipedia*. https://en.wikipedia.org/wiki/Ray_Kurzweil. (accessed March 8, 2016).
Renault-Nissan to introduce ten self-driving vehicles by 2020. In *The Washington Times*. http://www.washingtontimes.com/news/2016/jan/7/renault-nissan-to-introduce-10-self-driving-vehicl/. (accessed March 15, 2016).
Rich, Elaine, Kevin Knight, and Shivashankar B. Nair. *Artificial Intelligence*. New Delhi, India: Tata McGraw Hill Education Private LImited, 2009. http://itechnocrates.weebly.com/uploads/5/5/2/7/55270269/artificial_intelligence.pdf. (accessed March 12, 2016).
Rochefoucauld, Le Duc de La. 1665. *Réflexions ou Sentences et Maximes Morales de La Rochefoucauld (Dernier Texte)*. London, England: Louys Glady, 1880.
———. 1665. *Maxims of Le Duc de La Rochefoucauld*. Translated by John Heard. Boston, MA: International Pocket Library, 1917.
Russell, Stuart, and Max Tegmark. 2015. Think-tank dismisses leading AI researchers as Luddites. In *Future of Life Institute*. http://futureoflife.org/2015/12/24/think-tank-dismisses-leading-ai-researchers-as-luddites/. (accessed March 5, 2016).
Salamon, Anna. 2015. Why startup founders have mood swings (and why they may have uses). In *Less Wrong*. http://lesswrong.com/lw/n0v/why_startup_founders_have_mood_swings_and_why/. (accessed January 22, 2016).
Shoham, Yoav, and M. Tennenholtz. "On the synthesis of useful social laws for artificial agent societies." Presented at the Proceedings of the Tenth National Conference on Artificial Intelligence, San Jose, CA 1992, 276-81.
Stoffer, Shawn. 2000. Cyc: Building HAL. In *Department of Computer Science, University of New Mexico*. http://www.cs.unm.edu/~storm/docs/Cyc.htm. (accessed March 15, 2016).
Tucker, Patrick. 2014. The Military's New Year's Resolution for Artificial Intelligence. In *Defense One*. http://www.defenseone.com/technology/2014/12/militarys-new-years-resolution-artificial-intelligence/102102/?oref=search_Roadmap for AI. (accessed March 15, 2016).
———. 2015. The Pentagon is nervous about Russian and Chinese killer robots. In *Defense One*. http://www.defenseone.com/threats/2015/12/pentagon-nervous-about-russian-and-chinese-killer-robots/124465/?oref=d-river. (accessed March 15, 2016).
Uszok, Andrzej, Jeffrey M. Bradshaw, James Lott, Matthew Johnson, Maggie Breedy, Michael Vignati, Keith Whittaker, Kim Jakubowski, and Jeffrey Bowcock. "Toward a Flexible Ontology-Based Policy Approach for Network Operations Using the KAoS Framework."

Presented at the The 2011 Military Communications Conference (MILCOM 2011) 2011, 1108-14.
Valéry, Paul. 1937. "Our destiny and literature." In *Reflections on the World Today*. Translated by Francis Scarfe, 131-55. New York City, NY: Pantheon Books, 1948.
Webster, George. 2011. The future of airport security: Thermal lie-detectors and cloned sniffer dogs (25 November 2011). In *CNN*. http://edition.cnn.com/2011/11/25/tech/innovation/future-airport-security/index.html - ! (accessed March 15, 2016).
Weld, Daniel, and Oren Etzioni. "The first law of robotics: A call to arms." Presented at the Proceedings of the National Conference on Artificial Intelligence (AAAI 94) 1994, 1042-47.
Wellington, Naomi. "Whole brain emulation: Invasive vs. non-invasive methods." In *Intelligence Unbound: The Future of Uploaded and Machine Minds,* edited by Russell Blackford and Damien Broderick, 178-92. Chichester, West Sussex, England: John Wiley, 2014.
Work, Robert O., and Shawn Brimley. 2014. 20YY: Preparing for War in the Robotic Age. In *Center for a New American Security Publications*, Center for a New American Security. http://www.cnas.org/20YY-Preparing-War-in-Robotic-Age - .Vuhfh8c4Vqk. (accessed March 15, 2016).

---

1 With apologies to P. Valéry, *Our Destiny*; A. C. Clarke, Future Isn't (article); S. Jobs, Future Isn't, and others.

2 With apologies to D. McDermott, Artificial intelligence meets natural stupidity.

3 Yogi Berra: https://upload.wikimedia.org/wikipedia/commons/f/f1/1953_Bowman_Yogi_Berra.jpg (21 January 2016); J. Golden Kimball: https://rsc.byu.edu/sites/default/files/j%20golden%20kimball.jpg; Arthur C. Clarke: http://www.famousauthors.org/famous-authors/arthur-c-clarke.jpg; Steve Jobs: http://images.huffingtonpost.com/2015-07-17-1437155465-1424656-SteveJobs.jpg; Cassette: http://lifelibertytech.com/wp-content/uploads/2012/10/Talk-by-Steven-Jobs-Cassette.jpg

4 A version of the saying that uses the word "ain't" is often credited to Yogi Berra: "The future ain't what it used to be." According to Future Is Not, Quote Investigator, Future Is Not, Quote Investigator: "The baseball great Yogi Berra writing in his 1998 volume *The Yogi Book* did claim that he used this expression. A precise timeframe was not given, but the saying was accompanied with a picture from 1974. Yogi also offered an interpretation: 'I just meant that times are different. Not necessarily better or worse. Just different.'"

5 As is the case with many prominent figures, some of the stories told about Elder Kimball never happened. "Once a nephew came to J. Golden and said: 'Well, Uncle Golden, I heard another J. Golden yarn today.' 'I'll bet the damn thing isn't genuine,' wheezed the old man. 'Seems like all the stories told these days are either about me or Mae West" (A. Fife *et al.*, *Saints*, p. 314). Cf. E. A. Eliason, *J. Golden Kimball Stories*, p. 67.

6 See the full title of Y. Berra, *Yogi Book*.

7 See Future Isn't (video), Future Isn't (video); A. C. Clarke, Future Isn't (article). Clarke was prescient in anticipating many future developments.

8 S. Jobs, Future Isn't. See Y. Heisler, In 1983 Speech for many examples of what Jobs foretold accurately.

9 For a popular summary of initial evidence of some of the complex physiological, social, and cultural changes that technology is working upon us, see, e.g., S. Greenfield, *Mind Change*.

10 A. C. Clarke, Future Isn't (article), p. 4.

11 Paul Valéry: http://www.katakrak.net/sites/default/files/events/paul_valery.jpg (see also https://quoadsubjectum.files.wordpress.com/2012/08/paulvalery3.jpg, http://www.aphorism4all.com/images/1355814906.jpg)

12 The biography of Valéry on the website of "The Poetry Foundation" summarizes (Paul Valéry (1871-1945), Paul Valéry (1871-1945)):

> Paul Valéry occupies a position in the history of French letters that is at once strategic and highly problematic. Critics have affixed to him various labels, all of them partially correct. He has been called the last French symbolist, the first post-symbolist, a masterful classical prosodist, an advocate of logical positivism, and a cerebral narcissist. … [H]e is understood as having broken away from symbolism, as having rejected the cult of poetry for its own sake in favor of a cult of the mind. These views need not be contradictory. …
>
> Some facts about Valéry might predict a less than faultless comportment on Valéry's part during World War II and France's occupation by Germany: first, he had been
>
> quietly but strongly "anti-Dreyfusard" during the famous Dreyfus affair … Furthermore, Valéry was also friendly with Marshal Philippe Petain, one of the leaders of France's pro-German Vichy government.
>
> However, the poet did prove sympathetic to the Free French Movement led by General Charles de Gaulle, and of the Nazis he wrote in "War Economy for the Mind": "As for our enemies, we, and the whole world, know that their politics with regard to the mind has been reduced or limited for ten years to repressing the developments of intelligence, to depreciating the value of pure research, to taking often atrocious measures against those who consecrated themselves to these things, to favoring, even as far as endowed chairs and laboratories, worshippers of the idol to the detriment of independent creators of spiritual richness, and they have imposed on the arts as on the sciences the utilitarian ends which a power founded on declamations and terror pursues."

13 Future Is Not, Quote Investigator found one earlier occurrence of this saying than Valéry's:

> The earliest evidence of this saying located by [*Quote Investigator*] was published in 1937 in a journal called "Epilogue" within an article titled "From a Private Correspondence on Reality" by Laura Riding and Robert Graves. The authors who were both prominent literary figures asserted that the perception of the future had changed:
>
> > The human mind has reached the end of temporal progress: the future is not what it used to be, and people talk with less and less progenitive self-precipitation into the future, and behave with more and more fatally decisive immediacy. The future, that is, contains nothing but scientific development. It is an involuntary spending and manipulation of physical forces, empty of consciousness: it no longer matters.

In 1950, Mordecai M. Kaplan wrote: "Men say the future isn't what it used to be. Neither is the past. Both are in need of reconstruction, if we are to have a livable present" (M. M. Kaplan, Random Thoughts, cited in C. C. Doyle et al., *The Dictionary of Modern Proverbs*, p. 90).

14 P. Valéry, *Our Destiny*, pp. 135, 143-144.

15 Ibid., p. 142.

16 Luke 20:30-31.

17 L. D. d. L. Rochefoucauld, *Maxims*, 93, p. 29. L. D. d. L. Rochefoucauld, *Maxims*, 93, p. 21: « Les vieillards aiment à donner de bons préceptes, pour se consoler de n'être plus en état de donner de mauvais exemples. »

18 D&C 88:118; 109:7, 14.

19 http://i.huffpost.com/gen/3052140/images/o-PEELING-ORANGE-facebook.jpg

20 Image licensed from www.shutterstock.com. Image reference 67038073

21 See www.ihmc.us/groups/jbradshaw/.

22 For a summary of some of the unique aspects of IHMC's approach to the DARPA Robotic Challenge, written for the general reader, see M. Johnson et al., Seven Cardinal Virtues. For a video presentation that includes a description of the application of the principles of coactive design to the DARPA Robotic Challenge, see J. M. Bradshaw et al., Lessons Learned. For a detailed description of coactive design, see M. Johnson et al., Coactive Design.

23 H. Neubauer, *Curious Moments*, p. 634. Image licensed from Black Star / Alamy Stock Photo, Image Reference A2Y241.

24 Ibid., p. 634.

25 From A. Salamon, Why Startup Founders.

26 Many of the thoughts in this section are drawn from and paraphrased from ibid.

27 I don't know who first came up with this version of the saying, but it was not the Duc de La Rochefoucauld, to whom so many places on the Web attribute it falsely. I like this particular wording of the sentiment, which came from a talk I heard Elder Maxwell give in 1978 (N. A. Maxwell, The stern but sweet seventh commandment (Devotional, Salt Lake Institute of Religion, 8 December 1978), p. 8).

The original source from which the basic thought is derived is probably Thomas Huxley, in his Presidential Address to the British Association for the Advancement of Science, Liverpool Meeting, 14 September 1870 (T. H. Huxley et al., *Scientific Memoirs* 3, p. 580; T. H. Huxley, Biogenesis and Abiogenesis, p. 244): "But the great tragedy of Science — the slaying of a beautiful hypothesis by an ugly fact — which is so constantly being enacted under the eyes of philosophers, was played, almost immediately, for the benefit of Buffon and Needham."

28 As Anna Solomon described it (A. Salamon, Why Startup Founders):

> It's easy to point to the value in euphoria and optimism. You get lots of code written, recruit lots of funding and talent, write a perfect draft — it's the part of the cycle where you're drawn to working seventy hour weeks, checking off each and every item from your to-do list. But the "down" parts often feel like they're pointless at best, and dangerous or counterproductive at worst. … In our own pasts, we found ourselves wondering why our brains couldn't just hang on to the momentum — why they insisted on taking us through stupid detours of despair or shame before returning us back to apparent "forward motion."

29 Ibid.

30 Ibid.

31 Gartner Hype Cycle, Gartner Hype Cycle.

32 Ibid.

33 Gartner's 2015 Hype Cycle, Gartner's 2015 Hype Cycle.

34 How It Works, How It Works.

35 Ibid.

36 Cf. L. D. d. L. Rochefoucauld, *Maxims*, 106, p. 31: “Knowledge is an absolute mastery of details, and since of details there is no end, ours is always an imperfect and superficial knowledge.” L. D. d. L. Rochefoucauld, *Maxims*, 106, p. 24: « Pour bien savoir les choses, il en faut savoir le détail, et comme il est presque infini, nos connoissances sont toujours superficielles et imparfaites. »

To their credit, Google has invested heavily in empirical research on driving in real-world environments and has openly critiqued failures of autonomous capabilities in these situations. For example, see this assessment, briefly summarized in R. Lindner, Google Warning:

> Since 2009, when Google began working on self-driving cars, its fleet has so far covered more than two million kilometers in autonomous mode. During that time, there were, according to the group's report, 272 cases where a failure of the autonomous capabilities had been determined. In 69 cases, the driver had intervened to avert accidents. Subsequent simulations had shown the probability that in 13 of these cases an accident would have occurred had there been no intervention by a human driver. In a blog post Google pointed out, however, that the number of such incidents is declining. Of the 13 dangerous situations, eight had occurred eight in 2014, while only five took place last year. But Google also warned that the number could rise again if the autonomous cars were used under complex conditions — for example in bad weather.

37 http://johnnyholland.org/wp-content/uploads/DonaldANorman.jpg.

38 Bradshaw to Chair; New Scientific Advisory Council.

39 D. A. Norman, *Design of Future Things*, p. 13.

40 R. Lindner, Google Warning.

41 W. C. Fields in an unknown movie clip.

42 D. A. Norman, *Design of Future Things*, pp. 113, 116.

43 Ibid., p. 116.

44 M. Harris, Will Nissan Beat Google and Uber.

45 See, e.g., R. Lindner, Google Warning:

> Meanwhile, Google is looking for ways for its autonomous vehicles to close ranks with the auto industry. John Krafcik, who is in charge of the project, said on Tuesday at a press conference that took place near the auto show in Detroit that Google wants to enter into collaborations with many different companies. A number of automakers have sought contact with Google. A few weeks ago there was speculation that Google held discussions with the American automaker Ford shortly before the announcement of the alliance. Some wonder whether Ford could be contracted to manufacture Google’s next generation of self-driving cars. So far, these speculations have not materialized. Ford’s CEO Mark Fields distanced himself last week at the electronics show in Las Vegas from such cooperation and said he would not be limited to the role of a supplier for the technology industry.

46 M. Harris, Will Nissan Beat Google and Uber. See also the discussion of some of IHMC's recent work with Nissan on intelligent fleet management services in Come Hear Jeff Bradshaw.

47 Jeffrey M. Bradshaw, 6 January 2016, Image Reference IMG_6117.jpg.

48 AP Photo/Terry Chea in Renault-Nissan to Introduce.

49 J. Markoff, For Now Self-Driving Cars.

50 N. Bostrom, *Superintelligence*.

51 Doug Lenat's Cyc can be seen as a philosophical grandfather to such efforts. In 1999, Lenat wrote: "HAL was a general artificial intelligence, and Cyc is the closest thing that exists in the world to that kind of general artificial intelligence" (S. Moody, Brain). See also D. B. Lenat, From 2001.

52 https://upload.wikimedia.org/wikipedia/commons/thumb/f/f6/HAL9000.svg/1024px-HAL9000.svg.png.

53 A. C. Clarke, Foreword, p. xi.

54 http://www.cyc.com/wp-content/uploads/2015/04/kbase.png.

55 For brief summaries of some of the most common criticisms of Cyc, see Cyc; S. Stoffer, Cyc. For Lenat's own assessments of his work on Cyc, see R. V. Guha et al., Cyc: A midterm report; D. B. Lenat et al., CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks; D. B. Lenat et al., *Building Large Knowledge-based Systems*; D. B. Lenat, From 2001. For a recent video perspective by Lenat, see D. B. Lenat, Computers with Common Sense.

56 https://commons.wikimedia.org/wiki/File:IBM_Watson.PNG.

57 Anonymous, How IBM.

58 https://en.wikipedia.org/wiki/File:Raymond_Kurzweil_Fantastic_Voyage.jpg.

59 Cited in S. Armstrong et al., Who Knows?, p. 55.

60 Ibid., p. 55. For more detail, see S. Armstrong, Assessing Kurzweil.

61 John Rennie, quoted in Ray Kurzweil . See also S. Armstrong, Assessing Kurzweil.

62 S. Armstrong et al., Errors, Figure 1. Cf. S. Armstrong et al., Who Knows?, p. 51 Figure 3.1.

63 S. Armstrong et al., Who Knows?, p. 50.

64 Ibid., p. 51 Figure 3.1.

65 S. Armstrong et al., Errors, Figure 2. Cf. S. Armstrong et al., Who Knows?, p. 52 Figure 3.2.

66 S. Armstrong et al., Who Knows?, pp. 50-51.

67 R. Kurzweil, *Singularity*.

68 J. Gray, *Immortalization Commission*, pp. 214.

69 Ibid., p. 214.

70 Ibid., p. 216.

71 Ibid., p. 216. For an example of a recent title that embraces "Christian transhumanism," see R. Cole-Turner, *Transhumanism and Transcendence*.

72 http://www.fhi.ox.ac.uk/brain-emulation-roadmap-report.pdf. Brief, informal critiques of this roadmap can be found in the blogosphere include http://blog.ciphergoth.org/blog/2010/02/24/doug-clow-whole-brain-emulation-roadmap/ ; http://lproven.livejournal.com/279235.html ; http://blog.ciphergoth.org/blog/2010/02/20/david-matthewman-whole-brain-emulation-roadmap/ .

73 N. Bostrom, *Superintelligence*, p. 31. For examples of some of the preliminary thinking going on in this arena, see R. A. Koene, Feasible Mind Uploading; N. Wellington, Whole Brain Emulation; N. Bostrom, *Superintelligence*, pp. 30-36.

74 C. C. Green et al., *Emerging Cognitive Neuroscience*, p. 80.

75 For a readable summary of the state of the art in the study of the nervous system and its functions, see F. H. Gage, What Is the Brain Good For?.

76 P. Tucker, Pentagon Is Nervous.

77 For the terms of reference of this study, see F. Kendall, Terms of Reference. For additional background on this study, see P. Tucker, Military's New Year's Resolution. For remarks by the Deputy Defense Secretary, Robert O. Work, that quote from a draft of the study, see P. Tucker, Pentagon Is Nervous. For a brief overview of some of the "myths" of autonomy for the general reader, see J. M. Bradshaw et al., Seven Deadly Myths. For a video presentation for a general academic audience describing and illustrating these myths, see J. M. Bradshaw, Human-Agent-Robot Teamwork.

78 For early efforts to explore computational approaches for these laws, see, e.g., R. Clarke, Asimov's laws of robotics: Implications for information technology, Parts 1 and 2; Y. Shoham et al., On the synthesis of useful social laws for artificial agent societies; D. Pynadath et al., Revisiting Asimov's first law: A response to the call to arms; D. Weld et al., The first law of robotics: A call to arms.

79 See, e.g., J. M. Bradshaw et al., Policy-Based Governance; A. Uszok et al., Toward a Flexible Ontology-Based Policy Approach for Network Operations Using the KAoS Framework; J. M. Bradshaw et al., KAoS; S. Joseph et al., Digital Policy Management.

80 R. O. Work et al., 20YY.

81 See C. C. Green et al., *Emerging Cognitive Neuroscience*, p. 95: "While modeling the whole brain is highly unlikely in the next two decades, it is not unreasonable to imagine that significant subsystems could be modeled. Moreover, it seems likely that increasingly sophisticated cognitive systems will be constructed in those two decades that, while not aiming to mimic processes in the brain, could nonetheless perform similar tasks well enough to be useful, especially in constrained situations."

82 G. Webster, Future of Airport Security. According to the CNN article in which this image appeared, the thermal imaging system behind this image portends a new approach to detecting deception visually:

> Feeling guilty? Got something to hide? A team of UK-based researchers claim to have developed a thermal lie-detection camera that can automatically spot a burning conscience.
>
> The system could be used during customs interviews and at passport control to check whether people entering the country are giving a true account of themselves.
>
> The thermal-imaging camera captures variations in facial temperature in response to questioning. "When someone is making something up on the spot, brain activity usually changes and you can detect this through the thermal camera," said professor Hassan Ugail, who leads the research.

At present, the UK's Home Office and HM Revenue & Customs are sponsoring the system's development, but will not reveal the name of the airport where it's being tested.

83 C. C. Green et al., Emerging Cognitive Neuroscience, pp. 18-41. The study, which was published in 2008, was specifically looking out two decades, i.e., to the period ending in 2028.

84 See, e.g., J. M. Bradshaw et al., Sol; L. Bunch et al., Human-Agent Teamwork; L. Bunch et al., Principles for HCI Interaction Design 2; R. M. Jones et al., Modeling and Integrating. For a readable summary of early efforts for the general reader, see New Tools .

85 L. Bunch et al., Principles for HCI Interaction Design 2.

86 L. Bunch et al., Human-Agent Teamwork; J. M. Bradshaw et al., Multi-Agent Systems. For general overviews of software agent technology, see J. M. Bradshaw et al., Human-Agent Interaction; J. M. Bradshaw, *Software Agents*.

87 See, e.g., a summary of the view of most mainstream AI researchers expressed by Thomas G. Dietterich in CACM Oct. 2015 Video.

88 E.g.:

- Allen Institute for Artificial Intelligence (https://en.wikipedia.org/wiki/Allen_Institute_for_Artificial_Intelligence, http://allenai.org)
- Centre for the Study of Existential Risk (https://en.wikipedia.org/wiki/Centre_for_the_Study_of_Existential_Risk, http://cser.org)
- Future of Humanity Institute (https://en.wikipedia.org/wiki/Future_of_Humanity_Institute, http://www.fhi.ox.ac.uk)
- Future of Life Institute (https://en.wikipedia.org/wiki/Future_of_Life_Institute, http://thefutureoflife.org)
- Global Catastrophic Risk Institute (https://en.wikipedia.org/wiki/Global_Catastrophic_Risk_Institute, http://gcrinstitute.org)
- Institute for Ethics and Emerging Technologies (https://en.wikipedia.org/wiki/Institute_for_Ethics_and_Emerging_Technologies, http://ieet.org)
- Machine Intelligence Research Institute (MIRI), formerly the Singularity Institute for Artificial Intelligence (SIAI) (https://en.wikipedia.org/wiki/Machine_Intelligence_Research_Institute, https://intelligence.org)
- OpenAI (https://en.wikipedia.org/wiki/OpenAI, https://www.openai.com/blog/introducing-openai/)

89 E.g., AAAI Presidential Panel; AI100.

90 E.g., T. G. Dietterich et al., Rise of Concerns.

91 E.g., N. Bostrom et al., *Global Catastrophic Risks*; N. Bostrom, Superintelligence.

92 E.g., LessWrong (https://en.wikipedia.org/wiki/LessWrong; http://lesswrong.com).

93 Open Letter: Research Priorities; Autonomous Weapons: Open Letter.

94 R. D. Atkinson, 2015 ITIF Luddite Award Nominees.

95 S. Russell et al., Think-Tank Dismisses.

96 The result of this assimilation has sometimes been called the "AI effect," which "occurs when onlookers discount the behavior of an artificial intelligence program by arguing that it is not *real* intelligence" (AI Effect). This phenomenon was famously lamented by Douglas Hofstadter (D. R. Hofstadter, *Gödel, Escher, Bach*, p. 601):

> It is interesting that nowadays, practically no one feels that sense of awe any longer — even when computers perform operations that are incredibly more sophisticated than those which sent thrills down spines in the early days. The once-exciting phrase "Giant Electronic Brain" remains only as a sort of "camp" cliché, a ridiculous vestige of the era of Flash Gordon and Buck Rogers. It is a bit sad that we become blasé so quickly.
>
> There is a related "Theorem" about progress in AI: once some mental function is programmed, people soon cease to consider it as an essential ingredient of "real thinking." The ineluctable core of intelligence is always in that next thing which hasn't yet been programmed. This "Theorem" was first proposed to me by Larry Tesler, so I call it Tesler's Theorem: *"AI is whatever hasn't been done yet."*

The problem was characterized by the well known AI pioneer Marvin Minsky as a sort of argument by "redefinition" against AI: i.e., an effort to minimize any appearance of progressin the field of AI "by continually modifying the definition of intelligence in order to exclude all artificially reproduced phenomena" (R. Cordeschi, Discovery of the Artificial, p. 233. Cordeschi cites M. Minsky, Steps, p. 396, but I have been uable to track down anything in this paper or in other writings by Minsky that corresponds to this idea.)

Part of the problem is the fact that some popular definitions make humans the measure of AI research progress, e.g.:

> Artificial intelligence is the science of making machines do things that would require intelligence if done by men" (M. Minsky, *Semantic Information Processing,* p. *v*).
>
> Artificial Intelligences (AI) is the study of how to make computers do things which, at the moment, people do better (E. Rich et al., *Artificial Intelligence*, p. 3.

97 G. B. Hinckley, Thanks, p. 88.

98 http://mi.byu.edu/wp-content/uploads/2013/10/Nibley-1.jpg.

99 B. J. Petersen, *Nibley*, pp. 97-98.

100 D&C 89:4.