

The Future Isn't What It Used to Be:1 Artificial Intelligence Meets Natural Stupidity2

Part 6: Autonomous Weapons and Natural Stupidity

Editor's Note: The following is Part 6 in a series that expands upon a presentation given at the Second Interpreter Science and Mormonism Symposium: Body, Brain, Mind, and Spirit at Utah Valley University in Orem, Utah, 12 March 2016. A book based on the first symposium, held in 2013, has been published as Bailey, David H., Jeffrey M. Bradshaw, John H. Lewis, Gregory L. Smith, and Michael L. Stark. Science and Mormonism: Cosmos, Earth, and Man. Orem and Salt Lake City, UT: The Interpreter Foundation and Eborn Books, 2016. For more information, including free videos of these events, see <http://www.mormoninterpreter.com>.

The fifth article in this series can be found at <https://ldsmag.com/how-close-are-we-to-being-able-to-model-the-brain/>

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.'"

Figure 1 about here3

One area of technology that is fraught with momentous consequences is the proliferation of autonomous weapons — in other words, weapons that have significant capabilities for performing sensing, reasoning, and decision-making on their own. This was one of many topics that I explored in depth as a member of the 2015 Defense Science Board Summer Study on Autonomy.⁴ While I won't comment on the details of the study nor try to summarize the committee's consensus, I want to share my personal thoughts and concerns about the future development and deployment of weapons that are capable of operating more or less on their own. Such weapons, available in both the cyber and the physical domains, will be increasingly used 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.



****Figure 2 about here****

My interest in this topic has grown over the last fifteen years as our research group at the Institute for Human and Machine Cognition (IHMC) has worked on technological solutions to the problem of policy-based governance of intelligent systems, with a long-term vision that embraces the spirit of Isaac Asimov's laws of robotics.⁵ We call our digital policy services framework KAOs.⁶ Significant efforts are underway at the world's largest tech companies "to create a standard of ethics around the creation of artificial intelligence."⁷



Figure 3 about here8

Artificial Intelligence as a Pillar of Modern Military Strategy

Over the past fifty years, much of the significant progress in Artificial Intelligence (AI) has been due to funding from the United States Department of Defense (DoD). The resulting developments in AI can be grouped into three major waves:⁹

1. “The first wave (1950–1970) launched the academic field of computer science, opened an era of discovery and set the foundation for signal processing, computer vision, computer speech and language understanding.
2. The second wave (1970–1990) saw codification of knowledge in expert systems, using rule bases, and beginnings of simple machine inference to do reasoning (think things like computer chess), along with exploration of computer architectures, specialized for AI applications.
3. The third wave (1990–present) launched the era of large scale robotics, including autonomous machines, along with real breakthroughs in the use of neural network architectures, inspired by better understanding of how the brain works.”



U.S. DEPARTMENT OF DEFENSE

HOME ▾ NEWS ▾ ARTICLE

Deputy Secretary: Third Offset Strategy Bolsters America's Military Deterrence

By Cheryl Pellerin
DoD News, Defense Media Activity

PRINT | E-MAIL



WASHINGTON, Oct. 31, 2016 — The Pentagon's [Third Offset Strategy](#) pursues next-generation technologies and concepts to assure U.S. military superiority, but the real focus is strengthening U.S. conventional deterrence to make sure wars don't happen, Deputy Defense Secretary Bob Work says.



Deputy Defense Secretary Bob Work speaks at the Center for Strategic and International Studies about the Pentagon's third offset strategy in Washington, D.C., Oct. 28, 2016. DoD photo by Navy Petty Officer 1st Class Tim D. Godbee

He spoke last week at a Center for Strategic and International Studies' event titled, Assessing the Third Offset Strategy. Joining Work during the opening plenary was [Air Force Gen. Paul J. Selva](#), vice chairman of the Joint Chiefs of Staff.

At the forum, defense leaders discussed DoD's drive to identify innovative capabilities that will ensure U.S. military superiority over what the deputy secretary calls "pacing competitors."

Pacing Competitors

Work said the third offset begins by focusing on competitors who are developing advanced

****Figure 4 about here****¹⁰

These developments are now seen as so important to the DoD that they have been labeled its "third offset strategy."¹¹ Like the Cold War strategies of nuclear deterrence (first offset) and general technological superiority in the face of decreasing manpower (second offset), it is hoped that a third offset based on AI would provide a decisive advantage for the US military in future confrontations.¹²

Amazon 'Reviewing' Its Website After It Suggested Bomb-Making Items

By AMIE TSANG SEPT. 20, 2017



A British TV station said Amazon's website was prompting customers to buy all the ingredients to make a bomb. Richard Drew/Associated Press

RELATED COVERAGE



Google and Facebook Face Criticism Over Ads Targeting Racist Sentiments
SEPT. 15, 2017



Amazon Profits Fall Far Short of Expectations, but Investors Shrug
JULY 27, 2017



BITS
Tech Roundup: A Defense of 'Pro-Ads'
APRIL 6, 2017

LONDON — Amazon said on Wednesday that it was reviewing its website after a British television report said the online retail giant's algorithms were automatically suggesting bomb-making ingredients

****Figure 5 about here****¹³

What Are Some of the Challenges?

Several issues complicate the implementation of the “third offset” by the US military, including the fact that research and development of AI is now generously funded by (and largely controlled by) private companies rather than the DoD and the fact that the United States no longer holds a monopoly on significant scientific advances in the field.¹⁴

Unlike the nuclear weapons first deployed in World War II, the proliferation of autonomous weapons would not be constrained by the difficulties of a given nation's ability in performing sophisticated refinement of rare elements. Rather it 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. In principle, such capabilities could be developed in and sold from anyone's garage — so long as that garage has a good Internet connection.

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 custom tailored to specific purposes.¹⁵ Like the combination of bomb-making parts that, until recently, were cheerfully suggested by Amazon's recommendation algorithms to anyone who asked the right questions,¹⁶ AI

algorithms and code that are “good enough” to include in advanced weaponry are widely available everywhere.

Add all this to the fact that “weapons” are no longer confined to specialized military hardware or even conventional computers, but can reside and proliferate in the billions of connected gadgets of all kinds in our homes, workplaces, and public sites. Security for such devices is a daunting prospect.¹⁷ Billions of such devices are already in use and will easily dwarf the numbers of traditional computing devices in the coming few years. According to the Defense Science Board report: “This immense, sparsely populated space of interconnected devices could serve as a globe-spanning, multi-sensing surveillance system or as a platform for massively proliferated, distributed cyber-attacks — or as an immense test range for real-world, non-permissive testing of large-scale autonomous systems and swarms.”¹⁸

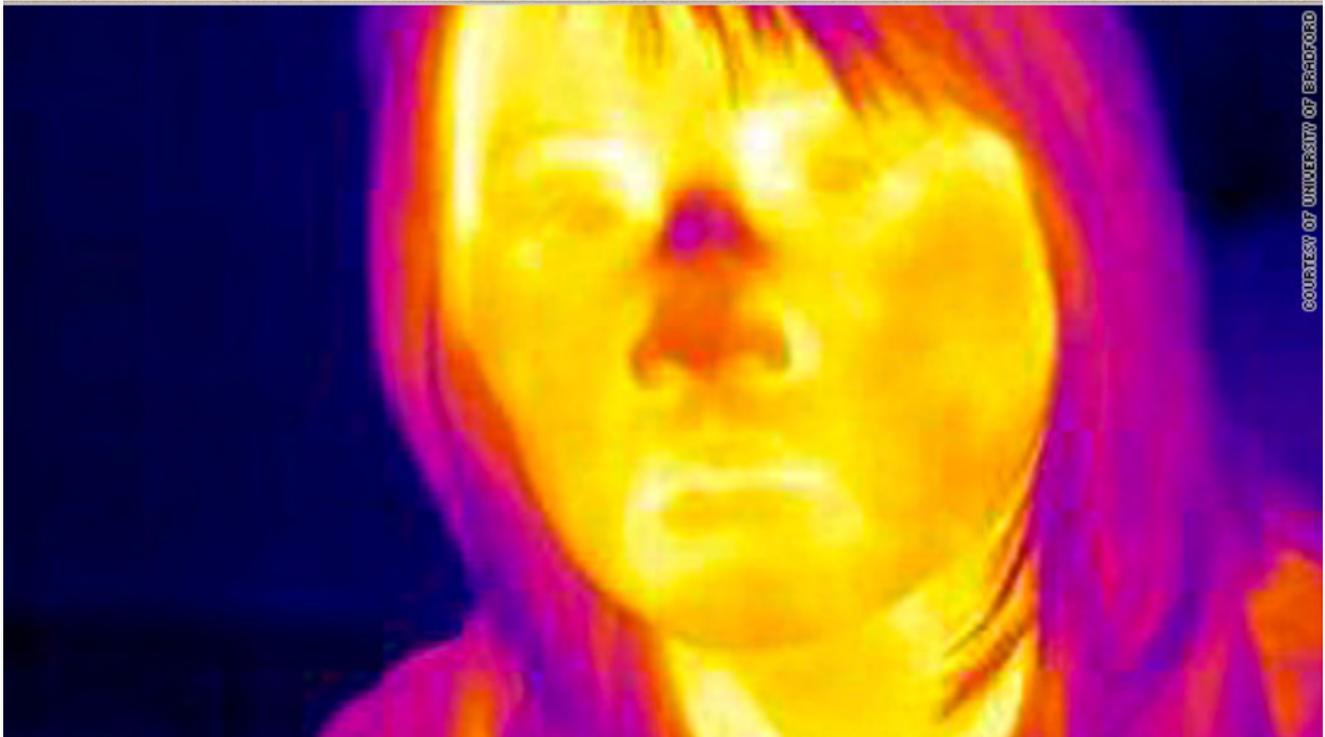


Figure 6 about here¹⁹

In previous articles in this series, we have given examples of the overblown expectations of scientific researchers about the near-term future of AI. Just to prove that others besides researchers can entertain wild speculations, at the initial meeting of a National Academies study some years ago, our group was told that one of the questions one sponsor had asked us 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.²⁰

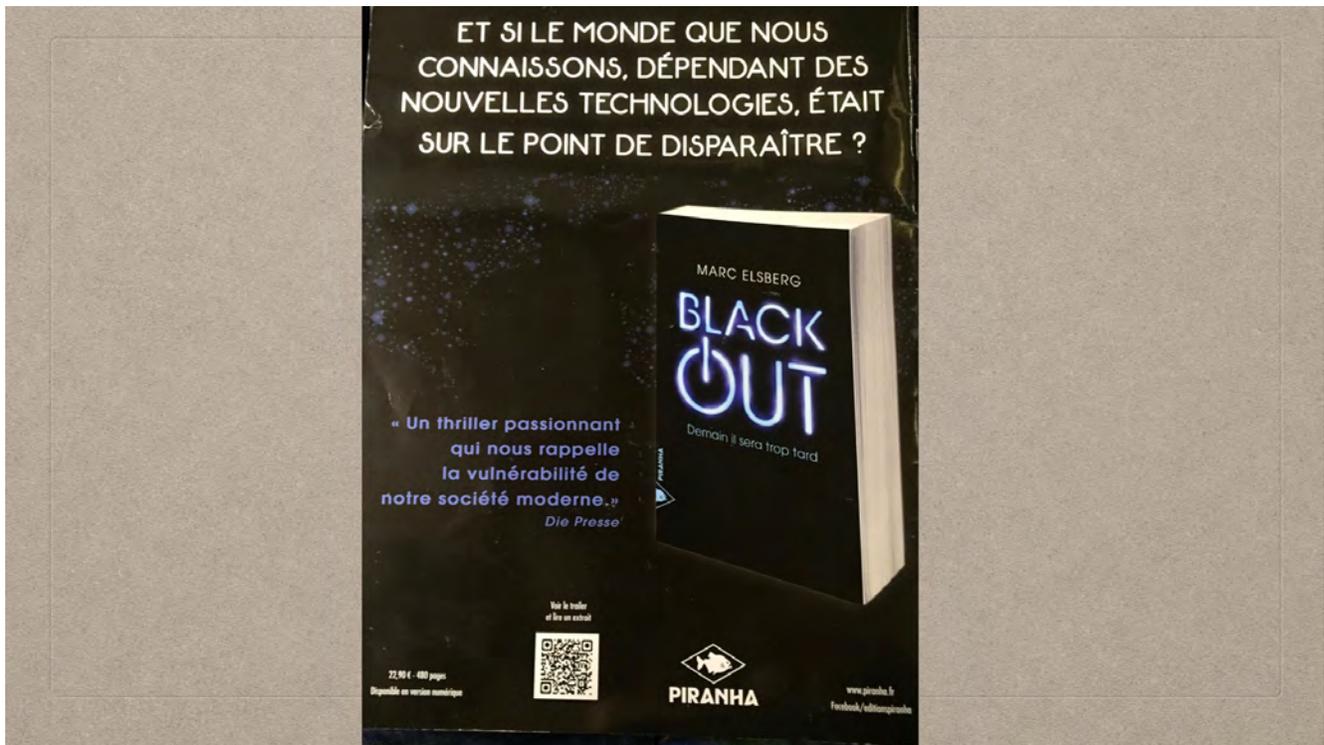


Figure 7 about here

The Rise of Cyber Warfare

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. For this reason, the DoD has elevated cyber security as a “national priority” and has established well-funded organizations to carry out its missions, such as the “US Cyber Command.”²¹

The motivation for cyber warfare waged against nations, organizations, and individuals is not merely political but is also economic. There is a flourishing worldwide “underground economy” that exploits the money to be made in “cybercrime, money laundering, and information security” breaches.²² Groups with a “motivation to find exploitable defects in widely used [software] ... are willing to pay anyone who can find and exploit these weaknesses top dollar to hand them over, and never speak a word to the companies whose programmers inadvertently wrote them into software in the first place.”²³ Far from the ideals of the Internet pioneers who imagined open access to information across all borders, we are facing the future of a “splinternet” fragmented by geopolitics and commercial interests.²⁴

Following hard on the heels of the enormous destructive power of two major hurricanes, damaging wildfires, and an 8.1 magnitude earthquake in Mexico, was the news of the September 7, 2017 theft of detailed personal and financial information at Equifax. This cyber disaster affected the lives and credit of up to 143 million people in the United States.²⁵ It has been called “one of the gravest breaches in history,”²⁶ but it is barely a drop in the bucket in the sea of information already available online about individuals. It provides a small foretaste of what portend to be greater confusions and disruptions of people’s private and public lives ahead.²⁷

Consider not only individual mavericks who manipulate online information for personal profit or political ends, but more importantly the increasing number of well-financed and carefully targeted efforts to create misinformation, invent false identities, and disrupt critical infrastructure with the goal of “wreak[ing] havoc all around the Internet — and in real-life American communities.”²⁸ For example, as early as 2008 the DoD publicly disclosed information “from multiple regions outside the United States, of cyber intrusions into utilities, followed by extortion demands. ... We have information that cyberattacks have been used to disrupt power equipment in several regions outside the United States.”²⁹

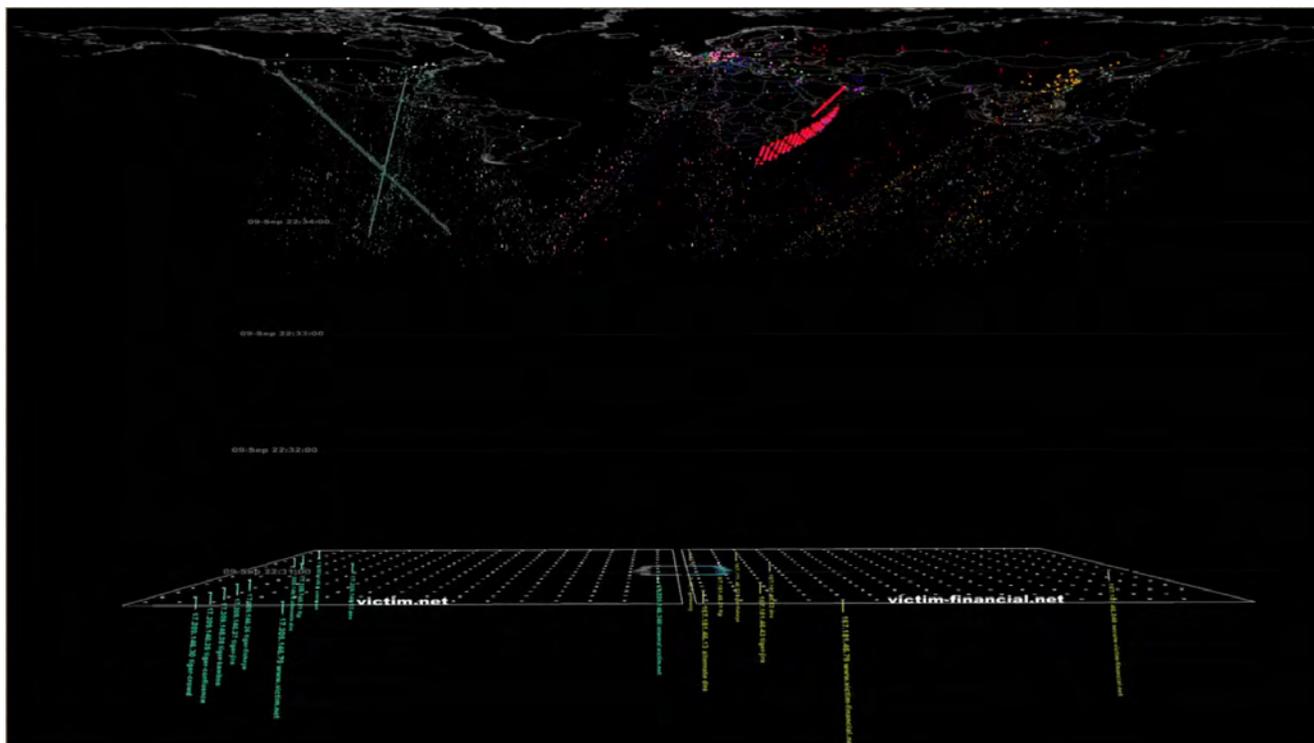


Figure 8 about here

IHMC’s Sol cyber framework, here using simulated data, shows one of the approaches our research team developed in response to a government request to address the (impossible) challenge of visualizing and interacting with the entire Internet in real time so as to make sense of whatever important events were going on at the moment.³⁰ We have had a “live,” real-time version of such a display continuously working on IHMC’s own network for some years now. As you watch the “live” display, the graphics make it easy to see continuous waves of attacks from around the world attempt to penetrate our relatively obscure and unimportant website.

The patented design of this and similar IHMC-developed displays exploit specific, subtle properties of human perception and cognition, allowing large numbers of interesting events to pop out and be assimilated by the ambient vision system.³¹ In the image, 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.



Figure 9 about here

The Future of Artificial Intelligence

Our design philosophy for Sol was consistent with the emphasis of our research group on creating systems that enable human-agent-robot teamwork (HART) rather than developing Artificial Intelligence capabilities that are meant to work more or less on their own. A good illustration of the more common way of thinking in the standalone AI approach can be found in the work of Alan Turing. Turing, a famous early computer scientist, asked the question, “Can machines think?” He laid out an experiment in the form of a game.³² The challenger in the game is given the task of comparing the separate answers of a human and a machine in order to determine which is which.

By way of contrast to Turing’s game, our question has been “Can humans and machines think *together*?” The challenge in designing Sol was not to determine whether a machine could be so sophisticated that it could fool a human. Instead, Sol was designed as an early experiment in blurring the line between human and machine thinking – to understand what it might be like someday for humans and machines to be working together so closely and that it would seem as if the parties were thinking together.³³ To this end, the visual innovations of Sol were combined with software agents that were designed to collaborate with cyber analysts, working together to make sense of complex situations in rapid, real time.³⁴ Because cyber attacks can occur in microseconds, the responsibility for the most rapid kinds of reactions must be assigned to the agents while deliberative aspects of sensemaking and decision-making can benefit from a combination of human and machine abilities.

While mainstream researchers in Artificial Intelligence usually reject the prospects of an AI explosion, singularity, or apocalypse such as those popularized in the media,³⁵ they have been thinking more deeply of late about the future of AI. As a result of this thinking, there has been a recent proliferation of research institutions,³⁶ studies,³⁷ articles,³⁸ books,³⁹ blogs,⁴⁰ and open letters of concern⁴¹ 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 these researchers are sometimes painted to be,⁴² they are some of the top minds in the field, believers in the potential of AI for the good of humankind.⁴³



Figure 10 about here

Combating Natural Stupidity

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.⁴⁴ I share much of the optimism of President Gordon B. Hinckley who, like his predecessors, rejected unsound extrapolations of scripture and statements of Church leaders to justify apocalyptic panic in the face of natural disasters and technological advances.⁴⁵ He 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.⁴⁶

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.”



Figure 11 about here⁴⁷

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. However, 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:⁴⁸

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. “I am grateful to know,” wrote Truman G. Madsen, “that Jesus Christ suffered not only for our sins but for our stupid mistakes.”⁴⁹ And through the Atonement of Jesus Christ, declared Elder Jeffrey R. Holland, “we can escape the consequences of both sin and stupidity — our own or that of others — in whatever form they may come to us in the course of daily living.”⁵⁰ 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,”⁵¹ and, most important of all, rely on divine wisdom and grace to help 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 September 9, 2017).
- AI Effect. In *Wikipedia*. https://en.wikipedia.org/wiki/AI_effect. (accessed March 12, 2016).
- Alba, Davey. 2017. The world may be headed toward a fragmented “splinternet”. In *Wired*.
<https://www.wired.com/story/splinternet-global-court-rulings-google-facebook/>. (accessed September 22, 2017).
- Artificial Intelligence: What’s real, what’s not, and is this the DoD Third Offset? (Abstract of Panel Discussion). In *Tentative Agenda for the Defense Science Board Sixtieth Anniversary: Celebrating Innovation for National Security (20 September 2016)*.
<https://www.eiseverywhere.com/ehome/179363/408368/>. (accessed September 26, 2017).
- 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*.
<https://itif.org/publications/2015/12/21/2015-itif-luddite-award-nominees-worst-year-s-worst-innovation-killers>. (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).
- Bernard, Tara Siegel, Tiffany Hsu, and Ron Lieber. 2017. Equifax says cyberattack may have affected 143 million in the US. In *The New York Times*.
<https://www.nytimes.com/2017/09/07/business/equifax-cyberattack.html>. (accessed September 9, 2017).
- 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.jeffreybradshaw.net/publications/121102-Multi-Agent Systems Abstract v. 2.pdf](http://www.jeffreybradshaw.net/publications/121102-Multi-Agent%20Systems%20Abstract%20v.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://jeffreybradshaw.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.jeffreybradshaw.net/publications/CSIIRW KAoS paper-s.pdf](http://www.jeffreybradshaw.net/publications/CSIIRW%20KAoS%20paper-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 Niranjana Suri and G. Cabri, 259-84. Boca Raton, FL: CRC Press/Taylor and Francis, 2014. <http://www.jeffreybradshaw.net/publications/130529-PolicyDraft2May-jb-rm-jb.pdf>.
- Bradshaw, Jeffrey M. 2016. Designing Software Agents: Perils, Pitfalls, and Promise. Video Presentation in Two Parts. In *Vivint Innovation Center, Lehi, UT*. <http://www.jeffreybradshaw.net/>. (accessed September 9, 2017).
- . 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://jeffreybradshaw.net/videos/160108-Bradshaw-Human-Agent-Robot Teamwork Through Coactive Design-TU Delft-174th Dies Natalis Seminar-Robots-Tools or Teammates-edited.mp4](http://jeffreybradshaw.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).
- 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.jeffreybradshaw.net/publications/30-03-HCC.PDF>. (accessed March 15, 2016).
- Chen, Adrian. 2017. The Agency. In *The New York Times*.
<https://www.nytimes.com/2015/06/07/magazine/the-agency.html>. (accessed September 9, 2017).
- 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).
- 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.
- 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).
- 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 2015): 38-40. <http://cacm.acm.org/magazines/2015/10/192386-rise-of-concerns-about-ai/fulltext>. (accessed March 5, 2016).
- Eaglen, Mackenzie. 2016. What is the third offset strategy? In *Real Clear Defense*.
https://www.realcleardefense.com/articles/2016/02/16/what_is_the_third_offset_strategy_109034.html. (accessed September 26, 2017).
- Feigenbaum, Edward A., Jeffrey M. Bradshaw, Alexander Felfernig, Ali-Akbar Ghorbani, Sven Koenig, Sankar Pal, David M. W. Powers, Vijay Raghaven, Eunice Santos, Takahira Yamaguchi, and Yiyu Yao. "Panel on Top Ten Questions in Intelligent Informatics and Computing in Celebration of the Alan Turing Year." Presented at the World Intelligence Congress, Macau, China, 4-7 December 2012, 2012. <http://wi-consortium.org/blog/top10qi/-cfp>. (accessed 9 September 2017).
- Fountain, Henry. 2017. Apocalyptic thoughts amid nature's chaos? You could be forgiven. In *The New York Times*. <https://www.nytimes.com/2017/09/08/us/hurricane-irma-earthquake-fires.html>. (accessed September 9, 2017).
- 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).
- Greenberg, Gary. 2008. Hackers cut cities' power. In *Forbes*.
http://www.forbes.com/2008/01/18/cyber-attack-utilities-tech-intel-cx_ag_0118attack.html. (accessed May 17, 2016).
- Hacking power networks. In *Schneier on Security*.
https://www.schneier.com/blog/archives/2008/01/hacking_power_n.html. (accessed May 17, 2016).
- Hicks, Kathleen H., Andrew Hunter, Jesse Ellman, Lisa Samp, and Gabriel Coll. 2017. Assessing the Third Offset Strategy. In *CSIS: Center for Strategic and International Studies*. https://csis-prod.s3.amazonaws.com/s3fs-public/publication/170302_Ellman_ThirdOffsetStrategySummary_Web.pdf?EXO1GwjFU22_Bkd5A.nx.fJXTKRDKbVR. (accessed September 26, 2017).

- 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.
- Holland, Jeffrey R. "'Tomorrow the Lord will do wonders among you' [Joshua 3:5]." *Ensign* 46, May 2016, 124-27.
- 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).
- 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.jeffreybradshaw.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).
- 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).
- Leigher, William E. 2011. Learning to operate in cyberspace. In *US Naval Institute Proceedings Magazine, February 2011*, pp. 32-37. <https://www.usni.org/magazines/proceedings/2011-02/learning-operate-cyberspace>. (accessed September 26, 2017).
- Lohr, Steve. 2016. Stepping up security for an Internet of Things world (16 October 2016). In *The New York Times*. <http://www.nytimes.com/2016/10/17/technology/security-internet.html>. (accessed October 18, 2016).
- Madsen, Barnard N. *The Truman G. Madsen Story: A Life of Study and Faith*. Salt Lake City, UT: Deseret Book, 2016.
- Malcomson, Scott. *Splinternet: How Geopolitics and Commerce Are Fragmenting the World Wide Web*. New York City, NY: OR Books, 2016.
- Manjoo, Farhad. 2017. Serious, Equifax? This is a breach no one should get away with. In *The New York Times*. <https://www.nytimes.com/2017/09/08/technology/seriously-equifax-why-the-credit-agencys-breach-means-regulation-is-needed.html>. (accessed September 9, 2017).
- Markoff, John. 2016. As artificial intelligence evolves, so does its criminal potential (23 October 2016). In *The New York Times*. <https://www.nytimes.com/2016/10/24/technology/artificial-intelligence-evolves-with-its-criminal-potential.html>. (accessed October 25, 2016).
- . 2016. How tech giants are devising real ethics for Artificial Intelligence (1 September 2016). In *The New York Times*. <https://www.nytimes.com/2016/09/02/technology/artificial-intelligence-ethics.html>. (accessed September 2, 2016).
- 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.
- New tools counter cyber threats. In *IHMC Newsletter 10:1 (February 2013)*.
<http://www.jeffreybradshaw.net/publications/IHMCNewslettervol10iss1.pdf>.
(accessed March 15, 2016).
- Noessel, Christopher. *Designing Agentive Technology: AI That Works for People*. Brooklyn, NY: Rosenfeld Media, 2017.
- Norman, Donald A. *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading, MA: Addison-Wesley, 1993.
- . *The Invisible Computer: Why Good Products Can Fail, the Personal Computer is So Complex, and Information Appliances Are the Solution*. Cambridge, MA: The MIT Press, 1998.
- Offset strategy. In *Wikipedia*. https://en.wikipedia.org/wiki/Offset_strategy. (accessed September 26, 2017).
- One Hundred Year Study on Artificial Intelligence (AI100). In *Stanford University*.
<https://ai100.stanford.edu/>. (accessed September 9, 2017).
- An open letter: Research priorities for robust and beneficial Artificial Intelligence. In *Future of Life Institute*. <http://futureoflife.org/ai-open-letter/>. (accessed March 5, 2016).
- Pellerin, Cheryl. 2016. Deputy Secretary: Third Offset Strategy Bolsters America's Military Deterrence. In *DoD News, US Department of Defense*.
<https://www.defense.gov/News/Article/Article/991434/deputy-secretary-third-offset-strategy-bolsters-americas-military-deterrence/>. (accessed September 26, 2017).
- 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.
- Report of the Defense Science Board Summer Study on Autonomy (June 2016). 2016. In, Defense Science Board, Office of the Under Secretary of Defense for Acquisition, Technology and Logistics. http://www.jeffreybradshaw.net/publications/DSB_Summer_Study_on_Autonomy_Final_Report_-_June_2016.pdf. (accessed 9 September, 2017).
- 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).
- Rosenberg, Matthew, and John Markoff. 2016. The Pentagon's 'terminator conundrum': Robots that could kill on their own. In *The New York Times*.
<http://www.nytimes.com/2016/10/26/us/pentagon-artificial-intelligence-terminator.html>. (accessed October 25, 2016).
- 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).
- 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.

- Software as weaponry in a computer-connected world. In *The New York Times* (7 June 2016). <http://www.nytimes.com/2016/06/09/technology/software-as-weaponry...ected-world.html>. (accessed June 10, 2016).
- Tsang, Amie. 2017. Amazon 'reviewing' its website after it suggested bomb-making items (20 September 2017). In *The New York Times*. <https://www.nytimes.com/2017/09/20/technology/uk-amazon-bomb.html>. (accessed September 26, 2017).
- 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).
- Turing, Alan M. "Computing machinery and intelligence." *Mind* 59, no. 236 (October) (1950): 433-60.
- Underground Economy 2013 Conference, Lyon 2-6 September 2013. In *Republic of Serbia, Administration for the Prevention of Money Laundering*. <http://www.apml.gov.rs/eng961/novost/Underground-economy-2013-Conference,-Lyon-2-%E2%80%93-6-September-2013.html>. (accessed September 26, 2017).
- 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.
- Work, Robert O., and Shawn Brimley. 2014. 20YY: Preparing for War in the Roboti 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).

Endnotes

-
- 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 P. Tucker, Pentagon Is Nervous.
- 4 Summer Study on Autonomy; 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.

5 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.

6 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.

7 J. Markoff, How Tech Giants.

8 R. O. Work *et al.*, 20YY.

9 What's Real, What's Not, What's Real, What's Not.

10 C. Pellerin, Deputy Secretary: Third Offset.

11 For readable introductions to the “third offset strategy,” see, e.g., K. H. Hicks *et al.*, Assessing; M. Eaglen, What Is the Third Offset Strategy; C. Pellerin, Deputy Secretary: Third Offset. For a brief sketch of what future technology may bring, see M. Rosenberg *et al.*, Pentagon's ‘Terminator Conundrum’.

12 Offset Strategy.

13 A. Tsang, Amazon ‘Reviewing’.

14 Software as Weaponry.

15 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.”

16 A. Tsang, Amazon ‘Reviewing’.

17 S. Lohr, Stepping Up Security.

18 Summer Study on Autonomy, Summer Study on Autonomy, p 88.

19 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.

20 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.

21 For an early snapshot of the DoD's view of the establishment of the US Cyber Command, see W. E. Leigher, *Learning to Operate*, p. 32.

22 *Underground Economy*.

23 *Software as Weaponry*.

24 S. Malcomson, *Splinternet*; D. Alba, *World May Be Headed*.

25 T. S. Bernard et al., *Equifax Says Cyberattack*.

26 F. Manjoo, *Seriously, Equifax?*

27 E.g., J. Markoff, *As Artificial Intelligence Evolves*.

28 A. Chen, *The Agency*.

29 E.g., G. Greenberg, *Hackers Cut; Hacking Power Networks*.

30 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*.

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

32 A. M. Turing, *Computing machinery and intelligence*. See also E. A. Feigenbaum et al., *Alan Turing Top Ten Panel*.

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

34 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*.

For a video introduction to software agents and design principles, see J. M. Bradshaw, *Designing Software Agents*.

For easy-to-read introductions to the topic, see C. Noessel, *Designing Agentive Technology: AI That Works for People*; D. A. Norman, *The Invisible Computer: Why Good Products Can Fail, the Personal Computer is So Complex, and Information Appliances Are the Solution*; D. A. Norman, *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*.

35 See, e.g., a summary of the view of most mainstream AI researchers in T. G. Dietterich et al., *Viewpoint*.

36 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/>)

37 E.g., AAI Presidential Panel, AAI Presidential Panel; AI100, AI100.

38 E.g., T. G. Dietterich *et al.*, Viewpoint.

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

40 E.g., LessWrong (<https://en.wikipedia.org/wiki/LessWrong>, <http://lesswrong.com>).

41 Open Letter: Research Priorities; Autonomous Weapons.

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

43 S. Russell *et al.*, Think-Tank Dismisses.

44 A few examples starting at the most basic level: object-oriented programming, semantic technologies (e.g., OWL and other related W3C standards), speech recognition, industrial robot motion planning and localization, facial recognition and a host of other vision processing algorithms used in photography, industrial robotics, security, and cinema.

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, 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 progress in 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 unable to track down anything in the referenced paper or in other writings by Minsky that corresponds to this idea.)

Part of the problem in properly characterizing AI is the lamentable tendency of some popular definitions to 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).

45 Of course, the tendency toward apocalyptic panic is not confined to a small segment of Church members. See, e.g., H. Fountain, *Apocalyptic Thoughts*.

46 G. B. Hinckley, *Thanks*, p. 88.

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

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

49 B. N. Madsen, *Truman G. Madsen*, p. 107.

50 J. R. Holland, *Tomorrow*, p. 127.

51 D&C 89:4.