

Open Source Artificial Intelligence (AI)

Alex Reichenbach

The development of full artificial intelligence could spell the end of the human race.

- Stephen Hawking (quoted by *BBC News*)¹

With artificial intelligence we are summoning the demon.

- Elon Musk (quoted in *The Guardian*)²

There are strong financial and moral arguments for whether or not software in general should be open source. However, I would like to discuss a specific application of software that should be considered a special case: Artificial Intelligence (AI). Experts in the field are concerned about the current and future implications of such technology, but their caution does not preclude their continuing to research in great depth the unique possibilities presented by AI.^{3,4}

There are as many approaches to dealing with AI as there are players in the field. Apple continues to make profound steps forward in the industry, but the company keeps its strides internal for a capitalist edge. Google's AI labs make most of their research public, but they keep their datasets private. OpenAI makes the majority of its code base public, including all its data.

Without incentives or regulations requiring companies to adhere to any common standards around AI, many companies (Apple, Affectiva, Google)^{5,6} keep their cutting edge research private. However, with all

¹ Rory Cellan-Jones, "Stephen Hawking warns artificial intelligence could end mankind," *BBC News*, last modified December 2, 2014, accessed January 16, 2018, <http://www.bbc.com/news/technology-30290540>.

² Samuel Gibbs, "Elon Musk: artificial intelligence is our biggest existential threat," *The Guardian*, last modified October 27, 2014, accessed January 16, 2018, <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat.-existential-threat>

³ "Elon Musk," *OpenAI*, accessed January 16, 2018, <https://openai.com/press/elon-musk/>.

⁴ RESEARCH PRIORITIES FOR ROBUST AND BENEFICIAL ARTIFICIAL INTELLIGENCE, accessed January 16, 2018, <https://futureoflife.org/ai-open-letter/>.

⁵ Will Knight, "Apple Just Gave a Private Demo of Its Latest Self-Driving Tech," *MIT Technology Review*, last modified December 10, 2017, accessed January 16, 2018, <https://www.technologyreview.com/s/609743/apple-researchers-just-gave-a-private-demo-of-their-latest-self-driving-tech/>.

⁶ Scott Ferguson, "Microsoft Establishes New AI Research Lab," *Enterprise Cloud Network*, last modified July 13, 2017, accessed January 16, 2018, http://www.enterprisecloudnews.com/author.asp?section_id=571&doc_id=734535.

the concerns posed by AI, as explored later, it may be impossible to regulate the industry in a meaningful manner without enforcing an open source policy.

AI is fundamentally just a piece of software, like a chess bot or an Operating System (OS). Even if one opposes mandated open sourcing software of this type generally, there are still strong arguments for open sourcing AI. AI is fundamentally different from other software applications for two main reasons: its opacity and its potential impact.

Unlike other software applications, the clear, logical flow of an AI application is not always readily apparent. Even leaders in the field sometimes struggle with why an AI might take a certain action, and this air of mystery doesn't dissipate even with close analysis in scientific papers. In *The Unreasonable Effectiveness of Recurrent Neural Networks*, one of the most famous studies of this type of machine learning architecture, the very first line is, "There's something magical about Recurrent Neural Networks (RNNs)."⁷

Why might a "magical" or, more critically, an opaque program prove problematic? After all, if the program works in the vast majority of cases, why might the remainder where it doesn't work be troublesome? The reason is that those untested cases might be easily exploited and provide a playground for hackers and security researchers. Recently, security researchers defeated a complex, neural-network-based authentication method by simply changing a single pixel in an image.⁸ Though a slew of papers followed suggesting methods to mitigate such an attack, none provided a complete solution.⁹ Even if designed without evil intent, AI is rife with unexpected behavior possibilities that can

⁷ Andrej Karpathy, "The Unreasonable Effectiveness of Recurrent Neural Networks," Andrej Karpathy blog, last modified May 21, 2015, accessed January 16, 2018, <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>.

⁸ Adam Geitgey, "Machine Learning is Fun Part 8: How to Intentionally Trick Neural Networks," Medium, last modified August 16, 2017, accessed January 16, 2018, <https://medium.com/@ageitgey/machine-learning-is-fun-part-8-how-to-intentionally-trick-neural-networks-b55da32b7196>.

⁹ Jiawei Su, Danilo Vasconcellos Vargas, and Sakurai Kouichi, *One pixel attack for fooling deep neural networks*, 1, October 24, 2017, accessed January 16, 2018, <https://arxiv.org/pdf/1710.08864.pdf>.

cause massive failures. During the trillion dollar 2010 Flash Crash, the Dow Jones average dropped 998.5 points in only thirty six minutes, the second largest intraday point swing in history.¹⁰ The major suspect behind the drop was a combination of manipulatory algorithms and autonomous, AI bot responses.¹¹ And, despite the Federal Trade Commission's (FTC) response, since then another monetary drop has occurred in bitcoin, which suddenly dropped 15% as a result of trading bot activity.^{12,13} Unlike a normal piece of software, AI programs can be impossible for the FTC to monitor or regulate effectively. If even AI experts and creators can't anticipate what an algorithm may do, how can a regulatory official be expected to do so?

AI is fundamentally different from other pieces of software. It has spread wide and far, from virtual assistants to authentication methods to security systems, yet we cannot guarantee the output. AI is naturally opaque, and, without releasing its origin source code, it becomes a "black box." Such a powerful "black box" creates an untenable situation with the potential for disastrous consequences.

Opaque and potentially biased mathematical models are remaking our lives—and neither the companies responsible for developing them nor the government is interested in addressing the problem.

-Will Knight, *MIT Technology Review*¹⁴

¹⁰ John Bates, "Post Flash Crash, Regulators Still Use Bicycles To Catch Ferraris," TRADERS, last modified April 24, 2015, accessed January 16, 2018, <http://www.tradersmagazine.com/news/technology/post-flash-crash-regulators-still-use-bicycles-to-catch-ferraris-113762-1.html?ET=tradersmagazine%3Ae4256762%3A1181926a%3A&st=email>.

¹¹ Bates, "Post Flash," TRADERS.

¹² US v. Sarao, 392 Tyler Bridegan & Dina Moussa (7th Cir. Apr. 2017). Accessed January 16, 2018. <https://www.georgetownlawtechreview.org/u-s-v-sarao-the-flash-crash-and-a-new-effort-to-prosecute-market-manipulation-and-deceptive-trading-practices/GLTR-04-2017/>.

¹³ Jon Russel, "Bitcoin, Ethereum and almost every other cryptocurrency is plunging," TechCrunch, last modified January 16, 2018, accessed January 16, 2018, <https://techcrunch.com/2018/01/16/bitcoin-crypto-crashed-hard-part-deux/>.

¹⁴ Will Knight, "Biased Algorithms Are Everywhere, and No One Seems to Care," MIT Technology Review, July 12, 2017, 1, accessed January 16, 2018, <https://www.technologyreview.com/s/608248/biased-algorithms-are-everywhere-and-no-one-seems-to-care/>.

Though the most publicized and sensational troubles arising from AI are those of existential risk, they remain distant.¹⁵ However, AI does have current and profound risks. It has pervaded common life at a faster rate than our awareness of its effects can keep up. Banks and insurance providers have used AI for years.¹⁶ Their goal, as corporations, is overwhelmingly to maximize profits. However, when approaching a dataset with that mindset, one often comes away from that dataset with a model that enforces the bias of the data. Let us take a specific, relevant example: the student loan crisis. The student loan crisis is highly racialized, “While just 6% of white borrowers who earned a bachelor’s degree defaulted on their loans, 14% of Latino borrowers did so, as did 23% of black borrowers,” and, “Overall, nearly half of black borrowers defaulted, more than double the rate for white borrowers.”¹⁷ If a bank were to train a model on this dataset, the resulting model would advocate giving loans to non-whites much less frequently than to whites, despite the fact that it may be illegal to do so. Without those loans, bias against those already disadvantaged groups would be further reinforced, leading to an even greater divide in the dataset, making the model more and more racist.¹⁸

Further, the bias may be so insidious that the bank may not even recognize the implicit bias of the model. A federal regulatory committee may be able to catch the large cases of such injustice, but it certainly could not uncover all of them. Proof is thin and money riding high. Some companies won’t adhere to the law. Some will passively avoid the question of AI bias, as in O’Neil’s research.¹⁹ Others will actively fight

¹⁵ Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford, UK: Oxford University Press, 2014)

¹⁶ Dean Sioukas, "Do You Fear the Future Of AI In The Banking World?," *Forbes*, accessed January 16, 2018, <https://www.forbes.com/sites/forbestechcouncil/2017/10/16/do-you-fear-the-future-of-ai-in-the-banking-world/#761d0aa47246>.

¹⁷ KAITLIN MULHERE, "There's a Massive Racial Gap in Student Loan Defaults, New Data Show," *Time*, last modified October 17, 2017, accessed January 16, 2018, <http://time.com/money/4986253/race-gap-student-loan-defaults-debt/>.

¹⁸ For an example of an AI learning racism, reference below. It describes Tay, a microsoft bot that learned to communicate via the community. Designed originally off a clean dataset, it started out well with pro humanity comments. Soon it degenerated into quotes such as “Hitler was right I hate the jews.” Vincent, James. “Twitter Taught Microsoft’s Friendly AI Chatbot to Be a Racist Asshole in Less than a Day.” *The Verge*, *The Verge*, 24 Mar. 2016, www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist.

¹⁹ “O’Neil previously worked as a professor at Barnard College in New York and a quantitative analyst at the company D. E. Shaw. She is now the head of Online Risk Consulting & Algorithmic Auditing, a company set up to help businesses identify and correct the biases in the algorithms they use. But O’Neil says even those who know

and avoid it, as did the defendants in an ACLU case investigating AI bias in the Medicaid program.²⁰

Without making their model creation source code open source, violators may continue to shirk civil rights laws with impunity as there may be no way to uncover and prosecute their non-blatant offenses.

Beyond cases of subtle biases created by AI, there are also occasions where AI has the potential to lead to tangible, serious, physical dangers.²¹ Driverless cars provide a good example. According to an *Atlantic* estimate, driverless cars could save about 29,447 lives, or nearly 300,000 lives in a decade.²² Currently allowed in Nevada, Florida, California, and Michigan, autonomous vehicles are a modern example of an AI making life or death decisions about humans. The only way to be sure that those cars are making the correct decisions, however, is if their origin code is open source. This would ensure that the highest level of safety were in place, help answer questions of liability, and promote clarity in decision-making.

Further, to create the safest autonomous vehicles possible and to gain public trust, the industry must go open source. If all AI cars were homogenous or compatible, then unexpected interactions, like those caused by the stock market bots discussed earlier, could be averted. Just as the internet is secure in large part because of its open source code, so must manufacturers of autonomous cars follow that lead.²³ We cannot ask a car to write a report as to why it made a decision that ended a person's life (perhaps because

their algorithms are at a risk of bias are more interested in the bottom line than in rooting out bias. "I'll be honest with you," she says. "I have no clients right now." Knight, "Biased Algorithms," 1.

²⁰ Jay Stanley, "Pitfalls of Artificial Intelligence Decisionmaking Highlighted In Idaho ACLU Case," ACLU, last modified June 2, 2017, accessed January 16, 2018, <https://www.aclu.org/blog/privacy-technology/pitfalls-artificial-intelligence-decisionmaking-highlighted-idaho-aclu-case>.

²¹ There have already been several crashes involving autonomous cars (Tesla and Mercedes in this example), and it is predicted that an autonomous car will kill its first human soon in the years to come. Robbins, Martin. "Statistically, Self-Driving Cars Are about to Kill Someone. What Happens next?" *The Guardian*, Guardian News and Media, 14 June 2016, www.theguardian.com/science/2016/jun/14/statistically-self-driving-cars-are-about-to-kill-someone-what-happens-next.

²² Adrienne Lafrance, "Self-Driving Cars Could Save 300,000 Lives Per Decade in America," *The Atlantic*, accessed January 16, 2018, <https://www.theatlantic.com/technology/archive/2015/09/self-driving-cars-could-save-300000-lives-per-decade-in-america/407956/>.

²³ Muhammad Farooq-i-Azam, "Role of Free and Open Source Software in Computer and Internet Security," Arxiv, accessed January 16, 2018, <https://arxiv.org/pdf/1610.00640.pdf>.

it mistook a person for a light pole), but we can try to understand its source and work together to prevent a similar occurrence in the future.

Another realm in which open sourcing for AI is particularly important is in the realm of medical predictions. Deep Patient is an AI medical predictor that has proven surprisingly effective at identifying the potential onset of psychiatric disorders, particularly schizophrenia, a disease which has been notoriously difficult to predict.²⁴ Despite experts analyzing the algorithm in great depth, they remain at a loss about how it is able to predict schizophrenia, though they cannot deny its effectiveness or great potential. Nevertheless, experts, like Google's head of AI, John Giannandrea, remain skeptical. "If someone is trying to sell you a black box system for medical decision support, and you don't know how it works or what data was used to train it, then I wouldn't trust it."²⁵ In other words, no AI model is perfect and, without access to its source code, which is the closest thing to reasoning we can get from it, we expose ourselves to great risks with potentially life-altering consequences.

While concerns about the current uses of AI are certainly pressing, the implications of future AI developments are so monumental that we must also focus diligently on them today.²⁶ Future Of Life Institute, an organization whose purpose-letter was signed by Elon Musk, Steve Wozniak, Geoffrey Hinton, and many other great technological thinkers,²⁷ pours millions of dollars into addressing questions of AI existentialism and safety.²⁸

²⁴ Will Knight, "The Dark Secret at the Heart of AI," MIT Technology Review, last modified April 11, 2017, accessed January 16, 2018, <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>.

²⁵ Will Knight, "Forget Killer Robots—Bias Is the Real AI Danger," MIT Technology Review, last modified October 3, 2017, accessed January 16, 2018, <https://www.technologyreview.com/s/608986/forget-killer-robotsbias-is-the-real-ai-danger/>.

²⁶ Bostrom, *Superintelligence: Paths*, 205.

²⁷ RESEARCH PRIORITIES FOR ROBUST AND BENEFICIAL ARTIFICIAL INTELLIGENCE.

²⁸ *FutureOfLife* reported \$3,291,262 in revenue according to their 990 tax return form. Tegmark, Max. "Return of Organization Exempt From Income Tax." FutureOfLife, 2016, futureoflife.org/wp-content/uploads/2017/05/2016_990_4website.pdf?x57718.

One issue the organization draws particular attention to is the use of AI for military purposes. Though it may seem ludicrous to require the military to open source its work, doing so may be necessary to ensure the protection of non-combatant civilians. The military's relentless pursuit of autonomous weaponry and software may prove reckless if not reviewed by peers. Some of the most secure protocols throughout history, mostly originating from the military, have been open source (ssh, ftp, rsa). When a case involving a military individual is publicly tried, we all learn why that soldier may have behaved a certain way. From that information, changes to training methods and protocol may be improved. However, we cannot conduct a public trial of a drone. If an autonomous drone were to bomb a civilian center, the public today would have no access to understanding why it did so. The closest we can get to understanding its action would be by releasing its source code. Only then might we learn from its mistakes and move forward to improve the accuracy and success of its future missions.

Further into the future, the possible AI atrocities only grow in abstractness and scale. Let us take a particularly non-intuitive fear of AI that is rarely addressed: mind crimes.²⁹ Consider a case where you benignly ask an AI to "Prepare the best cup of coffee I've ever had!" In response, an AI creates a complete copy of you inside a simulation. The simulation is precise down to your first thoughts, in order to increase accuracy. It determines that the best way to check whether you like a specific coffee is to give the simulated version of you the coffee. It runs thousands of simulations with different coffees, forcing the simulated version to drink each and measuring the amount of enjoyment produced. After the simulation is run, it removes the simulation. That simulated version of you only ever existed in the AI's memory, and yet it could think and be. It is still out for debate whether such a situation is even possible or whether it would necessarily be bad. Regardless, companies should be transparent if they are engaged in such controversial work. With companies not always honest and the stakes so high, the only way to keep companies honest and ethical is for their AI source code to be open. Developing such an open culture

²⁹ Bostrom, *Superintelligence: Paths*, 205.

around AI in its current infancy is critical, as it will set the precedent to be followed for decades to come, when the potential for AI risks will be even higher.

A major argument against mandatory open-sourcing is that it dulls a company's competitive edge, but companies can still file for patents to protect their work from competitors. They can also make their datasets proprietary, giving them a competitive advantage. Luckily, some companies are taking the lead toward open-sourcing. FaceBook, Google, and even Apple have released public research papers on their work into AI, a step in the right direction.³⁰ Baidu, a Chinese internet giant, has gone even further, making its entire autonomous driving platform open source.³¹

Despite the valid concerns raised by its continued growth, the potential upside of AI is too great to deter its increased development and use: lives saved by autonomous cars and medical algorithms; stable international financial systems from AI powered banking systems; and lives, young and old, bettered by automated assistance and even companionship. If implemented correctly, AI can be integrated into all our lives in an enriching, life-altering way. However, like atomic energy, AI is a tool rife with the potential for abuse. When used well, atomic energy can be one of the cleanest forms of energy, serving 11% of all households.³² When used poorly, it can destroy hundreds of thousands of lives.³³ Likewise, AI used well can make all our lives immeasurably better, but in the wrong hands it can unleash untold and unforeseeable harm. There is no single solution to address all the concerns raised by AI, but to open source it is certainly one of the best options available to us today.

³⁰ Kevin Lee and Serkan Piantino, "Facebook to open-source AI hardware design," FaceBook, last modified December 10, 2015, accessed January 17, 2018, <https://code.facebook.com/posts/1687861518126048/facebook-to-open-source-ai-hardware-design/>.

³¹ Roberto Baldwin, "Baidu updates its open-source autonomous driving platform," engadget, last modified October 26, 2017, accessed January 17, 2018, <https://www.engadget.com/2017/10/26/baidu-open-source-autonomous-driving-platform/>.

³² "World Statistics: Nuclear Energy Around the World," Nuclear Energy Institute, accessed January 17, 2018, <https://www.nei.org/Knowledge-Center/Nuclear-Statistics/World-Statistics>.

³³ "The Atomic Bombings of Hiroshima and Nagasaki," Atomic Archive, accessed January 17, 2018, http://www.atomicarchive.com/Docs/MED/med_chp10.shtml.

Bibliography

- “The Atomic Bombings of Hiroshima and Nagasaki.” Atomic Archive. Accessed January 17, 2018. http://www.atomicarchive.com/Docs/MED/med_chp10.shtml.
- Baldwin, Roberto. “Baidu updates its open-source autonomous driving platform.” engadget. Last modified October 26, 2017. Accessed January 17, 2018. <https://www.engadget.com/2017/10/26/baidu-open-source-autonomous-driving-platform/>.
- Bates, John. “Post Flash Crash, Regulators Still Use Bicycles To Catch Ferraris.” TRADERS. Last modified April 24, 2015. Accessed January 16, 2018. <http://www.tradersmagazine.com/news/technology/post-flash-crash-regulators-still-use-bicycles-to-catch-ferraris-113762-1.html?ET=tradersmagazine%3Ae4256762%3A1181926a%3A&st=email>.
- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. Oxford, UK: Oxford University Press, 2014.
- Cellan-Jones, Rory. “Stephen Hawking warns artificial intelligence could end mankind.” BBC News. Last modified December 2, 2014. Accessed January 16, 2018. <http://www.bbc.com/news/technology-30290540>.
- Eadicicco, Lisa. “Elon Musk is terrified that Google CEO Larry Page will accidentally create ‘evil’ robots that take over the world Read more at <https://www.businessinsider.com.au/elon-musk-terrified-larry-page-artificial-intelligence-2015-5#dgsbac5YL3pZMW7d.99>.” *Business Insider Australia*, May 25, 2015. Accessed January 11, 2018. <https://www.businessinsider.com.au/elon-musk-terrified-larry-page-artificial-intelligence-2015-5>.
- “Elon Musk.” OpenAI. Accessed January 16, 2018. <https://openai.com/press/elon-musk/>.
- Etzioni, Oren. “Perspective: Should Artificial Intelligence Be Regulated?” *Issues*. Accessed January 11, 2018. <http://issues.org/33-4/perspective-should-artificial-intelligence-be-regulated/>.
- Farooq-i-Azam, Muhammad. “Role of Free and Open Source Software in Computer and Internet Security.” Arxiv. Accessed January 16, 2018. <https://arxiv.org/pdf/1610.00640.pdf>.
- Ferguson, Scott. “Microsoft Establishes New AI Research Lab.” Enterprise Cloud Network. Last modified July 13, 2017. Accessed January 16, 2018. http://www.enterprisecloudnews.com/author.asp?section_id=571&doc_id=734535.
- Future of Life. Accessed January 9, 2018. <https://futureoflife.org/team/>.
- Geitgey, Adam. “Machine Learning is Fun Part 8: How to Intentionally Trick Neural Networks.” Medium. Last modified August 16, 2017. Accessed January 16, 2018. <https://medium.com/@ageitgey/machine-learning-is-fun-part-8-how-to-intentionally-trick-neural-networks-b55da32b7196>.

- Gibbs, Samuel. "Elon Musk: artificial intelligence is our biggest existential threat." *The Guardian*. Last modified October 27, 2014. Accessed January 16, 2018. <https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>.
- Hackett, Robert. "Alphabet's Eric Schmidt: 'I Was Proven Completely Wrong' About Artificial Intelligence." *Fortune*. Accessed January 11, 2018. <http://fortune.com/2017/02/15/eric-schmidt-rsa-artificial-intelligence/>.
- "Introducing OpenAI." *OpenAI Blog*. Entry posted December 11, 2015. Accessed January 8, 2018. <https://blog.openai.com/introducing-openai/>.
- Jiawei Su, Danilo Vasconcellos Vargas, and Sakurai Kouichi. *One pixel attack for fooling deep neural networks*. October 24, 2017. Accessed January 16, 2018. <https://arxiv.org/pdf/1710.08864.pdf>.
- Karpathy, Andrej. "The Unreasonable Effectiveness of Recurrent Neural Networks." Andrej Karpathy blog. Last modified May 21, 2015. Accessed January 16, 2018. <http://karpathy.github.io/2015/05/21/rnn-effectiveness/>.
- Knight, Will. "Apple Just Gave a Private Demo of Its Latest Self-Driving Tech." *MIT Technology Review*. Last modified December 10, 2017. Accessed January 16, 2018. <https://www.technologyreview.com/s/609743/apple-researchers-just-gave-a-private-demo-of-their-latest-self-driving-tech/>.
- . "Biased Algorithms Are Everywhere, and No One Seems to Care." *MIT Technology Review*, July 12, 2017. Accessed January 16, 2018. <https://www.technologyreview.com/s/608248/biased-algorithms-are-everywhere-and-no-one-seems-to-care/>.
- . "The Dark Secret at the Heart of AI." *MIT Technology Review*. Last modified April 11, 2017. Accessed January 16, 2018. <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>.
- . "Forget Killer Robots—Bias Is the Real AI Danger." *MIT Technology Review*. Last modified October 3, 2017. Accessed January 16, 2018. <https://www.technologyreview.com/s/608986/forget-killer-robots-bias-is-the-real-ai-danger/>.
- Lafrance, Adrienne. "Self-Driving Cars Could Save 300,000 Lives Per Decade in America." *The Atlantic*. Accessed January 16, 2018. <https://www.theatlantic.com/technology/archive/2015/09/self-driving-cars-could-save-300000-lives-per-decade-in-america/407956/>.
- Lee, Kevin, and Serkan Piantino. "Facebook to open-source AI hardware design." Facebook. Last modified December 10, 2015. Accessed January 17, 2018. <https://code.facebook.com/posts/1687861518126048/facebook-to-open-source-ai-hardware-design/>.

- MULHERE, KAITLIN. "There's a Massive Racial Gap in Student Loan Defaults, New Data Show." Time. Last modified October 17, 2017. Accessed January 16, 2018. <http://time.com/money/4986253/race-gap-student-loan-defaults-debt/>.
- Musk, Elon. "OpenAI Technical Goals." *OpenAI Blog*. Entry posted June 20, 2016. Accessed January 8, 2018. <https://blog.openai.com/openai-technical-goals/>.
- Quora. Accessed January 9, 2018. <https://www.quora.com/Should-artificial-intelligence-be-regulated>.
- RESEARCH PRIORITIES FOR ROBUST AND BENEFICIAL ARTIFICIAL INTELLIGENCE. Accessed January 16, 2018. <https://futureoflife.org/ai-open-letter/>.
- Russel, Jon. "Bitcoin, Ethereum and almost every other cryptocurrency is plunging." TechCrunch. Last modified January 16, 2018. Accessed January 16, 2018. <https://techcrunch.com/2018/01/16/bitcoin-crypto-crashed-hard-part-deux/>.
- Sioukas, Dean. "Do You Fear the Future Of AI In The Banking World?" Forbes. Accessed January 16, 2018. <https://www.forbes.com/sites/forbestechcouncil/2017/10/16/do-you-fear-the-future-of-ai-in-the-banking-world/#761d0aa47246>.
- Stanley, Jay. "Pitfalls of Artificial Intelligence Decisionmaking Highlighted In Idaho ACLU Case." ACLU. Last modified June 2, 2017. Accessed January 16, 2018. <https://www.aclu.org/blog/privacy-technology/pitfalls-artificial-intelligence-decisionmaking-highlighted-idaho-aclu-case>.
- US v. Sarao, 392 Tyler Bridegan & Dina Moussa (7th Cir. Apr. 2017). Accessed January 16, 2018. <https://www.georgetownlawtechreview.org/u-s-v-sarao-the-flash-crash-and-a-new-effort-to-prosecute-market-manipulation-and-deceptive-trading-practices/GLTR-04-2017/>.
- "World Statistics: Nuclear Energy Around the World." Nuclear Energy Institute. Accessed January 17, 2018. <https://www.nei.org/Knowledge-Center/Nuclear-Statistics/World-Statistics>.