Inherent Negatives of Artificial Intelligence

Laura Meléndez

Rutgers University
mmelendez.laura@gmail.com

Abstract

This paper analyses artificial intelligence (AI) technology from a business social responsibility perspective. It applies various corporate social responsibility frameworks, including Rubin and Carmichael's (2018) inherent negatives concept, stakeholder theory, and triple bottom line theory, to identify aspects of the AI technology that pose potential risks to societal order. The paper uses textual and context review and analysis of literature, studies and current AI news reports, as well as analysis of video content and speeches in the public domain, with focus on the United States (US). The inherent negatives identified are: jobs elimination due to the automation of mechanical tasks, dehumanization of work, algorithmic bias, insufficient data sets, lack of algorithm transparency and complex systems interdependence. The paper defines these inherent negatives, provides evidence of existence and explores business social responsibility implications, including the role of philanthropy. The discussion then turns to current governmental action that addresses the AI space, explaining the European Union's General Data Protection Regulation (GDRP) as well as other AI reports in the US and the United Kingdom. It also addresses some of the recently formed non-governmental organizations (NGOs), some born out of philanthropic efforts, that are working to balance AI's inherent negatives. Lastly, this paper outlines areas for future AI business social responsibility study.

Keywords: Algorithm, artificial intelligence, bias, data, GDPR, inherent negatives, job elimination, NGOs, social responsibility, stakeholders, triple bottom line.

Introduction

The Merriam-Webster dictionary defines artificial intelligence (AI) as the capability of a machine to imitate intelligent human behavior. Research company IDC predicts that by 2021 organizations will spend $52.2 billion annually on AI-related products. Economists and analysts believe businesses will realize billions in savings and gains from their AI investment (Vanian, 2018). In other words, AI is big business.

Talk about AI's potential can be traced back to the 1950s, with Alan Turing defining AI's ultimate realization as the moment when "a human communicating with a machine will not be able to distinguish the machine's response from a human's response" (Kile, 2012). By Turing's definition, AI became a reality on May 8th, 2018 when Google unveiled Google Duplex, a chat-box capable of maintaining a natural conversation in which humans were unable to recognize they were interacting with a robot (Ghaffary, 2018). This is a remarkable technological achievement. It is also a moment of reckoning for society. Are there social drawbacks to AI?

This paper analyzes AI technology from a business social responsibility perspective, applying Rubin and Carmichael (2018) inherent negatives framework to identify aspects of the technology that pose potential risks to societal order, the steps taken to address those risks thus far and outlines areas for future study. The paper uses textual and context review and analysis of literature, studies and current AI news reports, as well as analysis of video content and speeches in the public domain, with focus on the United States (US).

How does AI work?

During the 2017 Stanford Graduate School of Business talk Artificial Intelligence is the New Electricity, Andrew Ng explained that the AI deployed today mostly leverages supervised learning, using a large data set to understand and predict the response to a specific question, or "map an A to B response". For example, if a developer wants the AI system to recognize a dog, the system will be fed millions of dog images in order for the algorithm to identify as many dog characteristics as possible. Computer science professionals have been working on developing these capabilities for years. Recent advancements on the supervised learning branch of AI have been made possible by the convergence of the availability of large data sets, increased computing speeds and improved storage capabilities, which together enable deep learning. With deep learning, computer networks process enormous amounts of information, recognizing patterns quickly, and with less coaching from humans, which accelerates new product development (Vanian, 2018). In the dog example, the AI system's deep learning lower layers recognize simple things like outlines or color; higher layers recognize complex details like fur or eyes, and the topmost layer brings it all together to identify the picture as a dog (Knight, 2017).

Society has benefited tremendously from these advancements. Recommendation engines like Netflix and You Tube allow consumers to easily find content that fit their taste. Free online tools simplify the important, albeit tedious, task of creating citations for academic work. From traveling and mapping tools like Kayak and Waze, that allow consumers to save time and money, to financial products that monitor fraud, AI applications surround the consumer today and their presence will continue to increase (Gordon-Murnane, 2018).

Understanding the social impact of developing and using AI technology in business is critical because this technology is already extensively leveraged and it impacts a wide range of stakeholders. Stakeholders are typically defined by their interest in a corporation (Vogel, 2005). For the purpose of this analysis, a stakeholder is broadly defined as anyone who is affected by the mass deployment of AI technology, among them: consumers, businesses, suppliers, shareholders, employees, unions, governments, academic institutions, media outlets, non-governmental organizations (NGOs) and advocacy groups.

What are the inherent negatives of AI?

Like any other business, advancements in AI also have negative aspects that are part of the process, which if unchecked, can exacerbate already troublesome global societal trends such as economic inequality and poverty. Rubin and Carmichael (2018) define inherent negatives as "business risks derived from the negative stakeholder impact of latent attributes in a given business model". Using this definition, jobs elimination due to automation of mechanical tasks, dehumanization of work, algorithmic bias, insufficient data sets, lack of algorithm transparency and complex systems interdependence are inherent negatives that must be addressed when designing AI systems. The following section describes each concept in detail, offering definitions, examples and hypotheses on possible societal implications.

Jobs elimination due to automation of mechanical tasks

As previously explained, AI systems work by imitating human behavior based on a large number of examples. Ng (2017) articulated a rule of thumb for product managers regarding AI implementation: anything that a typical human can do with at most a second of thought can probably be automated with AI now or soon.

Thus far, this process has worked best in tasks that are highly repetitive. It is not surprising that industries like manufacturing have implemented AI systems to cut costs on repetitive processes such as assembly line production. In the US, the human toll of this productivity measure has been the loss of the moderately paid jobs that fueled the American middle class during the second half of the 20th century. As of 2016, five million manufacturing jobs were lost in the US, with 88% of that loss attributed to automation made possible by AI systems (Long, 2016).

The retail vertical is also undergoing an AI-fueled transformation. Walmart, the biggest non-government employer in 22 states (Desjardins, 2017), is experimenting with robot janitors and grocery pickers (Meyersohn, 2018). Amazon, US's second $1 trillion company (Bhattarai, 2018) and one of the top three retailers in the world, along with Walmart and Alibaba (O'Grady, 2018), has not only famously automated its warehouses using robots ("Amazon Warehouse Robots: Mind Blowing Video", 2016), but is also aggressively experimenting with completely self-serving brick and mortar stores (Amazon Go, n.d.).

On the farming and agriculture front, AI systems have been used in novel ways. Kile (2012) cites Cooper and Sigalla 1996 & Sigalla 2000 in examples of automation of the milking process, where a chip is implanted on a cow to monitor milk production and quality, allowing the cow to remain in a pasture until it needs to be milked. At the time of milking, the cow goes to a gate that opens automatically and proceeds to enter a completely automated milking facility. A quick internet search reveals dozens of AI solutions specifically designed for the milking industry (Appendix A).

While these measures increase productivity and consequently, shareholder value, they also have vast societal impact. Applied AI systems can displace humans partially or completely and reduce the availability of mid and low-wage job opportunities typically fulfilled by people without a college degree. The elimination of these jobs becomes a critical social issue when juxtaposed with the educational attainment US census report that found 68% of adults 25 years old or older do not have a college degree (U.S. Census Bureau, 2016). Educational attainment does not change much when examined by gender; 67% of females and 68% of males do not have a college degree (there was no data available for LGBTQ demographics). However, the analysis shows stark differences when examined across racial demographics: 46% of Asians, 67% of white, 78% of blacks and 85% of Hispanics do not have a college degree, which suggests AI-fueled jobs elimination may disproportionately impact already marginalized communities. This is perhaps the most salient inherent negative of AI today.

AI-fueled automation enabled by deep learning is expected to also impact white collar professions in the near-future. Management consulting firm McKinsey predicts that about 30% of the activities in 60% of all occupations could be automated (Manyika & Sneader, 2018). Increases in underemployment and unemployment can contribute to rising economic inequality across the globe, which in a doomsday scenario could lead to mass idleness and increases in social issues like substance abuse and criminality (Kile, 2012).

Doomsday scenarios are avoidable. The same McKinsey report points out that AI systems will create additional labor demand of between 555 million and 890 million jobs globally (Manyika & Sneader, 2018), including increased demand across verticals such as healthcare and infrastructure. They also point out that 8% to 9% of this demand will be in entirely new occupations. The study highlights the need to develop workforce differently, emphasizing a combination of technology with social, emotional and cognitive skills over physical and manual skills (Manyika & Sneader, 2018). Forbes Insights goes a step further, envisioning specific new jobs that will directly work on AI systems: Trainers, who will teach AI systems how to perform; Explainers, who will bridge the gap between technologists and business leaders by explaining complex algorithms to non-technical professionals; and Sustainers, who will ensure that AI systems are operating as designed ("What Are The New Jobs In A Human + Machine World?", 2018). Job creation predictions sound encouraging, but these predictions must be further scrutinized due to their association to AI developers. McKinsey is part of Partnership for AI, a think-tank founded by tech companies that develop and sell AI, and Forbes' article "What Are The New Jobs In A Human + Machine World" was paid by Intel AI (Appendix B).

To be successful implementing AI technology to the benefit of society, businesses, unions, governments, NGOs and other stakeholders will need to collaborate to reimagine and structure AI workforce development programs and employment expectations.

Dehumanization of work

Humans working alongside and collaborating with machines sounds like a science fiction movie, but it is already a reality for millions of workers around the world. AI systems enable people and machines to work together, changing the way work is performed and creating new business dynamics. ("What Are The New Jobs In A Human + Machine World?", 2018).

Managers are expected to balance humans and machines as productivity demands continue to increase. Recent evidence indicates that without the proper guardrails, management expectations of human labor could shift, possibly demanding humans behave like machines. Amazon, an AI pioneer, has been widely criticized for their working conditions and low wages in fulfillment centers (Amazon Employees Speak Out About Workplace Conditions | NBC Nightly News, 2018). In a harrowing exposé, The New York Times' podcast The Daily described how XPO Logistics, a warehousing vendor for some of the biggest retailers in the world, ignored a pattern of work-related miscarriages for women working at a Verizon warehouse in Memphis, TN - while demanding pregnant employees continue with duties that included heavy lifting - in order to achieve instant delivery production goals ("The Human Toll of Instant Delivery", 2018).

Management treating laborers like machines is not a new dynamic. From slavery to sweatshops, it seems like unrestrained, humans will abuse power. Reich (2007) argues that these abuses are the logical consequence of intensifying competition to give consumers and investors better deals in an environment where laws have not been enacted to protect citizens from the social consequences of a changing technology landscape. As long as the deals are legal and they satisfy consumers and investors, corporations will pursue them (Reich, 2007). Companies that develop or leverage AI anywhere in their supply chains to increase productivity should not only ensure that current labor laws are respected, but they should also create specific and prescriptive management guidelines that ensure humans continue to be treated with dignity within human/machine collaborative environments. Yet, in the absence of laws requiring companies to do so, will they? Consumers should try to understand the human consequences of cheaper prices and services like same day delivery and demand humane treatment for all workers. But will consumers sacrifice consumeristic convenience in the name of human dignity?

Algorithmic bias and insufficient data sets

Ng (2017) explained that developers use massive amounts of data in order to train AI systems. In other words, a human formulates a question, determines the correct answer to that question, and then teaches the algorithm how to arrive at the predetermined correct answer by using a lot of examples. This process works well when questions and answers are universally objective, like the result of a mathematical equation or the chemical composition of a substance. However, during the 2017 Ted Talk "The era of blind faith in big data must end", Cathy O'Neil pointed out that this seemingly democratic process in reality carries human cognitive bias when applied against real world challenges where judgement calls are necessary to reach an answer.

A cognitive bias is anything that skews how the brain processes information (SciShow, 2015). When people analyze information, their brains are comparing the new information to biases in order to form an opinion (SciShow, 2015). Everyone has biases, including the developers designing the algorithms and testing the results. Practical evidence strongly suggests algorithmic bias is another inherent negative of AI systems.

The AI development process has shown weaknesses in the form of algorithmic bias in several ways and for different reasons. First, the data itself may carry bias. Computer science already has a concept that can be applied to this situation: "Garbage In, Garbage Out.", which implies bad input will result in bad output (GIGO, n.d.). For algorithmic AI training this concept could be applied as "Bias In, Bias Out". Just as good programming practices dictate that functions should check for valid input before processing (GIGO, n.d.), good AI practices should require data sets to be reviewed for bias before deeming them valid for model training. This is not a simple proposition because developers may not recognize the bias in the data until the model is trained, and then only if the bias is so pronounced that it makes results obviously skewed. For example, in 2014 Amazon ended the development of a recruiting AI tool because the program learned to discriminate against women (Hamilton, 2018). The AI tool downgraded resumes containing the words "women's" and filtered out candidates who had attended two women-only colleges (Hamilton, 2018). It seems that the data used to train the AI system consisted of predominantly male resumes submitted to Amazon over a 10-year period. The system concluded that men were preferable (Hamilton, 2018). In this example, the developers could have reviewed the demographic composition of the data to ensure it was objective, but they would not have been able to correct the bias. Public records show that Amazon's workforce does have a male bias. Only 39% of the Amazon global workforce are women (Molla, 2018) and that is reduced to 27% when looking only at technical positions; therefore, both the data and the system's conclusion were correct, but biased - "Bias In, Bias Out". AI systems results' testing may benefit from extensive correlation vs. causation analysis to ensure that the algorithm does not assume causation when it identifies correlation. It is important to recognize that by shutting down the hiring tool project, Amazon signaled that it does not want to perpetuate the female hiring bias in the future.

Amazon's failed AI hiring tool example also points to a second data issue that directly affects algorithmic bias - insufficient data. Even if the company wants to correct for an existing bias, there is not enough data available to train the system to reach a different conclusion. Amazon has done some work trying to correct the insufficient data issue in a different space, food rotting, using oversampling (Vanian, 2018). Oversampling is a computer science technique where developers direct how the algorithm learns by assigning heavier statistical "weights" to underrepresented data (Vanian, 2018). While promising, this technique also carries human judgement calls that may introduce bias: who decides how much correction is necessary?

Judgement calls are a third issue with algorithmic bias and data. AI models have failed when answers to questions are a matter of opinion rather than fact and involve many judgement calls (O'Neil, 2017). For example, in 2016 Beauty.AI launched a contest where artificial intelligence would determine which faces most closely resembled "human beauty", from pictures submitted from around the world. Roughly 6,000 people from more than 100 countries submitted photos in the hope of being recognized as beautiful (Levin, 2016). The AI system returned results that filtered out dark skins. In this case the incoming data set was wide and representative of many different examples, but the model was trained much more narrowly using a western, white sample to establish standards of attractiveness. These standards may have been more closely aligned with the developers' perception of beauty, and did not include enough minorities, ignoring that different parts of the world have different beauty standards.

Representation is also an issue with data. Personal assistants like Apple's Siri, Google's Assistant and Amazon's Alexa use a mix of voice recognition, noise reduction, search and other machine learning systems to work. Ng (2017) explained that these systems have been thoroughly tested in English and Chinese Mandarin, which at first glance are the top two most spoken languages in the world. However, when language data is filtered by primary language, Spanish replaces English as the second most spoken language. What pushes English to the second place in total numbers is the immense amount of people that speak English as a second language (Most Widely Spoken Languages in the World, n.d.). Therefore, spoken English around the world is most likely to have some accentuation, instead of sounding neutral as it does on TV. It is reasonable to question whether the current voice recognition models have been trained with enough data to fully recognize the range of spoken English - a quick search of "Alexa fails" seems to indicate it has not (Appendix C). Furthermore, there are over 7,000 languages spoken in the world today, and around 2,000 of those have an average of 1,000 speakers. Will there ever be enough samples to train voice recognition systems on these languages? While "Alexa fails" are somewhat of a joke today when people ask AI-powered personal assistants to play a tune or describe the weather, the question of representation becomes much more important when, for example, biometric devices start incorporating voice recognition as a means of identification and access (Gordon-Murnane, 2018).

In another failed AI example, users quickly transformed Microsoft chatbot Tay from the "trained" AI personality of a teenage girl to sex-crazed neo-Nazi in less than a day (Metz, 2018). Developers did not foresee users giving hateful input and had not built any logic to ignore it. Tay is a useful example of what can go wrong when an AI model is unleashed to learn and adapt using real time input without guardrails.

Questions about judgement and representation on AI algorithms force an examination of the developers' demographic composition. A review of diversity in the US workforce for the top tech giants' technology positions - the positions that are likely responsible for creating and testing AI algorithms – indicates women are underrepresented: Amazon has the most women at 27%, followed by Apple at 23%, Google at 20%, Facebook and Microsoft at 19% (Molla, 2018). Is it any wonder that the Amazon's AI hiring tool concluded men are a better fit? The data is more disheartening when looking at race composition in the US tech workforce: at Amazon only 2.6% of tech positions are filled by black employees and 3.5% are filled with Latinx employees; at Apple 7% are black and 8% are Latinx; at Google and at Facebook 1% are black and 3% are Latinx; and at Microsoft 2.7% are black and 4.3% are Latinx (Molla, 2018). Lack of diversity in tech departments may be one of the reasons experiments like Beauty.AI suffered from bias-through-omission when choosing the data to train the AI model. These issues become even more pressing with increasingly sophisticated AI systems, such as startup Affectiva, which promises to "emotion-enable apps, devices and digital experiences, so these can sense and adapt to facial expressions of emotion". Given the ample evidence of algorithmic bias and the dramatic cultural differences on what qualify as appropriate when showing emotion among different groups in the US and around the world, training an algorithm to identify emotions effectively requires careful consideration and enormous amounts of data to guarantee proper representation.

Algorithms are also built with a business bias toward selling. Facebook's extensively documented blunders during the 2016 election show how its algorithm was designed to serve users articles that were popular with their friends, in some cases reinforcing bias and spreading misinformation, to increase time spent on site and exposure to paid advertising. While successful on increasing profitability, these algorithms alienated users into information bubbles whose effects are still being studied (Borchers, 2018).

In a piece on voice recognition devices for The Atlantic, Shulevitz (2018) points out that unlike an internet search where users are exposed to various answers to the same question, AI-powered personal assistants only give the user what the algorithm determines is the best answer. Taking into consideration the multiple biases the data could contain, from unconscious biases embedded in the data, to unrepresentative samples to bias towards selling, it is important to consider whose perspective of "best" consumers really are getting in this unique result. Companies must consider creating new frameworks to ensure algorithms are pressure-tested to eliminate or at least minimize bias.

Lack of algorithm transparency and complexity

AI algorithms are a competitive advantage and they are fiercely protected. Secrecy is important to stop competitors from quickly duplicating systems. In deep learning AI systems that continue to learn and make connections as data becomes available, even the engineers that built the initial algorithms are not able to explain exactly how the system evolves and determines results (Knight, 2017). This lack of algorithm transparency becomes an inherent negative, particularly when the system results are unexpected.

The average consumer is not able to explain, for example, how Alexa can create a music station with just a few instructions, or how Stitch Fix found the perfect jeans. In these trivial scenarios, suddenly getting a country song on a Latin rock station or receiving an ill-fitting blouse are nothing more than small glitches without major repercussions. However, deep learning AI systems are being actively used on everything from policing to healthcare applications, to banking. Understanding how the algorithm made a decision becomes critical in these scenarios (Knight, 2017). Rudin and Wagstaff (2014) argued that a predictive model cannot truly be useful unless a human understands it, regardless of how accurate it is. Current applications of AI, such as recommendations to treat illness are exposed to the very real possibility that the algorithm may not have taken into consideration a unique individual detail, highlighting the need for human understanding of AI reasoning.

Algorithmic secrecy is not the only complexity in the current and future implementation of AI systems. In order to create the breakthrough insights that the business demands, AI systems are highly interconnected, acquiring data from multiple sources at multiple times. These systems are supported by complicated hardware infrastructures, from microchips to data servers and rely on internet connectivity. This interdependence paired with the lack of transparency in the algorithm makes it very complicated to find answers when things go wrong, as evidenced by Uber's self-driving car accident that resulted in one fatality in May 2018 (O'Kane, 2018). Uber determined that the mistake was in the algorithm, which may have "decided" it did not need to take evasive action even though it saw the pedestrian (O'Kane, 2018). However, the mistake could have been in the hardware (did the sensors see the pedestrian?), the speed of reaction (did the system have enough time to take evasive action?) or at multiple other points of automatic decisions that are part of this system. In this example Uber determined what the mistake was, but it was not able to explain specifically why the mistake happened. Their best explanation was that the system *possibly* flagged the pedestrian detection as a "false positive". In a life or death situation, is society willing to accept *possibly* as the best explanation a company can offer?

Are AI's inherent negatives being addressed?

The good news about inherent negatives is that once negatives are identified, companies can work to minimize or eliminate harm. AI systems are no exception and some promising work is already underway.

Governments across the globe are actively defining governance models for AI, even though there is no consensus on the best path forward. An analysis of three 2016 (Obama administration) Office of Science and Technology Policy reports concludes that the US government's understanding of AI relies heavily on the liberal notion of the free market, where the vision of the government's role as a regulator is limited (Cath, Wachter, Mittelstadt, Taddeo, & Floridi, 2018). Nevertheless, these reports acknowledge some of the social impacts of AI, such as the likelihood of jobs elimination and the need for government to develop public policy to ensure that AI does not increase economic inequality (Cath, Wachter, Mittelstadt, Taddeo, & Floridi, 2018). The 2016 reports also recognize the need for open and unbiased data sets and the need for more diversity in algorithmic development, although they do not offer tangible steps to address the issues (Cath, Wachter, Mittelstadt, Taddeo, & Floridi, 2018). The 2018 White House Summit on Artificial Intelligence for American Industry report (Trump administration) indicated that the summit focused on AI research and development, workforce development, regulatory barriers to AI innovation, and sector-specific applications of AI. Notably missing from this agenda was explicit focus on ensuring that AI's benefit is extended to all society. It remains to be seen how the US government will enact policy, as effects of AI on society become more evident and prevalent.

Government in the European Union (EU) and the United Kingdom (UK) has been much more proactive than the US. The EU's 2016 Civil Law Rules on Robotics treats AI as an enabling technology to robotics, but it calls for specific action along several social fronts: implementation of employment forecasts to monitor job trends; refocusing educational goals to develop the workforce, especially women, with the necessary digital skills; consideration of a new tax to offset the revenue loss and societal negative effects that can potentially be caused by unemployment, underemployment and economic inequality; and consideration of obligatory disclosures of savings made in social security contributions due to automation (Cath, Wachter, Mittelstadt, Taddeo, & Floridi, 2018). The EU report also recommends the creation of a ''European Agency for Robotics and AI'', to monitor AI trends and envisioning its future impact, and with advising public players (Cath, Wachter, Mittelstadt, Taddeo, & Floridi, 2018). The UK's 2016 House of Commons' Science and Technology Committee report on AI calls for both reliance on existing regulation and for development of new regulatory frameworks. In addition, the report suggests the creation of an independent commission to organize public debate about AI challenges (Cath, Wachter, Mittelstadt, Taddeo, & Floridi, 2018).

Moving beyond policy discourse to policy implementation, the General Data Protection Regulation (GDPR) went into effect in the EU on May 25, 2018 (Kaminski, 2018). The GDPR is a set of rules on algorithmic accountability that imposes transparency, process, and oversight on the use of computer algorithms to make significant decisions about human beings (Kaminski, 2018). Practical policing of accurate implementation of GDPR rules remains a challenge for the EU government, but the law signals the government's willingness to act in favor of societal order (Kaminski, 2018).

In addition to governmental action, several NGOs have formed in the last few years to investigate, highlight and address some of AI's inherent negatives in different ways, among them: the Center for Humane Technology, led and funded by a team of ex-industry insiders; AI NOW Institute, an NYU AI research think-tank; Algorithmic Justice League, a collective with the mission to fight algorithmic bias; AI 4 ALL, a youth education effort led by a partnership of universities across the US; the Leverhulme Centre for the Future of Intelligence, a University of Cambridge AI research think-tank; and the Ethics and Governance of AI Initiative, an ethics-focused partnership among several NGOs. Several self-regulating industry organizations have predictably formed as well, among them: Future of Life Institute and Partnership on AI.

Conclusion

AI is an exploding, fast-moving discipline that has already irreversibly changed consumers' expectations on consumption. From delivery speed to personalization, AI has opened consumers' eyes to a world of new possibilities. Yet, as with all technology, progress comes at a cost. AI runs the risk of exacerbating some of the social issues that the United Nations' Sustainable Development Goals are working to address, such as economic inequality and poverty (About the Sustainable Development Goals, n.d.). Moreover, AI impacts are far-reaching, affecting a wide range of stakeholders at a global scale.

Handy (2003) argued that in a knowledge economy - one dependent on intellectual capital, rather than goods production - good business is a community with a purpose, not a piece of property. With the decline of manufacturing businesses and the rise of technology companies, the current US economy appears to have crossed the threshold from industrial to knowledge economy. The 2017 Cone Communications CSR Study found that 78% of consumers want companies to address important social justice issues, and 76% of consumers will refuse to purchase a company's products or services upon learning it supported an issue contrary to their beliefs. These findings indicate that current consumer attitudes support Handy's definition of good business. Handy (2003) also defined purpose as "not to make a profit, but to make a profit in order to do something better". Using Handy's framework, in order to be successful, technology companies designing and leveraging AI systems must invest time and effort deliberately defining the "something better" that will guide their business, help them craft appropriate strategies to address AI's inherent negatives and realize the technology's potential to the benefit of the economy and society.

International and Managerial Implications

This paper has outlined some of the inherent negatives of AI technology evident today from a business responsibility perspective, but much research is needed to fully understand the effects of these negatives on different stakeholder groups and to devise strategies on how to address them. Furthermore, this list is by no means exhaustive - AI impacts on the environment, human psychology, politics and many other sociological areas were not examined. Some of the most immediate areas in need of exploration for the business social responsibility practice are: Is it possible to build more inclusive algorithms by diversifying the developer workforce? How can the AI industry's impact on economic inequality be measured? How should workforce development programs and academic institutions evolve in order to accommodate the needs of both businesses and potential employees? Should AI system users be not only informed, but compensated for the data they provide? Do companies that leverage AI have a responsibility to maintain human dignity? Do consumers have a responsibility to understand the impact of consumption and convenience on an AI-driven society? What's AI's environmental impact?

Analyzing the AI industry through the Triple Bottom Line (People, Planet, Profit) corporate social responsibility lens, it seems that the industry has invested a considerable amount of resources maximizing profits for itself and across industries, benefiting the economy - the "Profit" dimension. However, the industry has not addressed the "People" and "Planet" dimensions. Rudin and Wagstaff (2014) observed that research efforts in AI are not distributed according to the needs of society. The implementation of the EU's GDPR signals that managers in the industry must begin to consider equal investments in "People", analyzing how operations impact not only employees' wellbeing, but society in the form of economic inequality and poverty around the globe. Moreover, even though current laws are not yet demanding this, leaders must also begin to consider investments in "Planet" to understand, address and minimize the operational impact of AI on the environment.

The AI industry is creating societal dynamics that demand holistic re-imagination of the way society works. Efforts to address societal issues must include input from fields ostensibly unrelated to technology, such as philosophy and the humanities. To be effective, this work will require evolved governance and stakeholder collaboration models (Aakhus, & Bzdak, 2015) that reach across disciplines and countries to place social issues at the center of discourse and solution-seeking efforts.

References

Aakhus, M., & Bzdak, M. (2015). Stakeholder engagement as communication design practice. Journal of Public Affairs, 15(2), 188–200. doi:10.1002/pa

About the Sustainable Development Goals - United Nations Sustainable Development. (n.d.). Retrieved from https://www.un.org/sustainabledevelopment/sustainable-development-goals/

Amazon Employees Speak Out About Workplace Conditions | NBC Nightly News. (2018, September 02). Retrieved from https://www.youtube.com/watch?v=tvdyxXhVNRE&feature=youtu.be

Amazon Go. (n.d.). Retrieved December 8, 2018, from https://www.amazon.com/b?ie=UTF8&node=16008589011

Amazon Warehouse Robots: Mind Blowing Video. (2016, July 26). Retrieved from https://www.youtube.com/watch?v=cLVCGEmkJs0&feature=youtu.be

Artificial Intelligence. (n.d.). Retrieved from https://www.merriam-webster.com/dictionary/artificial intelligence

Borchers, C. (2018, January 14). Facebook invites you to live in a bubble where you are always right. Retrieved from https://www.washingtonpost.com/news/the-fix/wp/2018/01/14/facebook-invites-you-to-live-in-a-bubble-where-you-are-always-right/?noredirect=on&utm_term=.ca0c0c79bbfa

Bhattarai, A. (2018, September 04). Amazon becomes the country's second $1 trillion company. Retrieved from https://www.washingtonpost.com/business/2018/09/04/amazon-becomes-countrys-second-trillion-company/?noredirect=on&utm_term=.643ece424614

Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., & Floridi, L. (2018). Artificial intelligence and the 'Good society': The US, EU, and UK approach. Science and Engineering Ethics, 24(2), 505-528. doi:10.1007/s11948-017-9901-7

Cone Communications CSR Study. (2017, May 17). Retrieved from http://www.conecomm.com/research-blog/2017-csr-study

Desjardins, J. (2017, November 20). The biggest employer in every US state. Retrieved from https://www.businessinsider.com/the-biggest-employer-in-every-us-state-2017-11

Ghaffary, S. (2018, May 08). Listen in as Google's AI assistant has phone conversations with real people. Retrieved December 8, 2018, from https://www.recode.net/2018/5/8/17332380/google-io-2018-ai-assistant-phone-conversations

GIGO. (n.d.). Retrieved from https://techterms.com/definition/gigo

Gordon-Murnane, L. (2018). Ethical, explainable artificial intelligence: Bias and principles. Online Searcher, 42(2), 22.

Hamilton, I. A. (2018, October 10). Amazon built an AI tool to hire people but had to shut it down because it was discriminating against women. Retrieved from https://www.businessinsider.com/amazon-built-ai-to-hire-people-discriminated-against-women-2018-10

Handy, C. (2003). What's a business for? In C.K. Prahalad & M. E. Porter (Eds.), Harvard Business Review on Corporate Social Responsibility (pp. 65-83). Cambridge: Harvard Business School Press.

Kaminski, M. (2018). The right to explanation, explained LawArXiv. doi:10.17605/OSF.IO/RGEUS

Kile, F. (2012). Artificial intelligence and society: A furtive transformation. Ai & Society, 28(1), 107-115. doi:10.1007/s00146-012-0396-0

Knight, W. (2017, May 12). There's a big problem with AI: Even its creators can't explain how it works. Retrieved December 8, 2018, from https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/

Levin, S. (2016, September 08). A beauty contest was judged by AI and the robots didn't like dark skin. Retrieved from https://www.theguardian.com/technology/2016/sep/08/artificial-intelligence-beauty-contest-doesnt-like-black-people

Long, H. (2016, March 29). U.S. has lost 5 million manufacturing jobs since 2000. Retrieved from https://money.cnn.com/2016/03/29/news/economy/us-manufacturing-jobs/

Manyika, J., & Sneader, K. (2018, June). AI, automation, and the future of work: Ten things to solve for. Retrieved from https://www.mckinsey.com/featured-insights/future-of-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for

Metz, R. (2018, March 27). Microsoft's neo-Nazi sexbot was a great lesson for makers of AI assistants. Retrieved from https://www.technologyreview.com/s/610634/microsofts-neo-nazi-sexbot-was-a-great-lesson-for-makers-of-ai-assistants/

Meyersohn, N. (2018, December 07). Walmart turns to robots and apps in stores. Retrieved from https://www.cnn.com/2018/12/07/business/walmart-robot-janitors-dotcom-store/index.html

Molla, R. (2018, April 11). How Facebook compares to other tech companies in diversity. Retrieved from https://www.recode.net/2018/4/11/17225574/facebook-tech-diversity-women

Most Widely Spoken Languages in the World. (n.d.). Retrieved from https://www.infoplease.com/arts-entertainment/writing-and-language/most-widely-spoken-languages-world

Ng, A. (2017, February 02). Andrew Ng: Artificial Intelligence is the New Electricity. Retrieved December 8, 2018, from https://www.youtube.com/watch?v=21EiKfQYZXc&feature=youtu.be

Stanford Graduate School of Business talk

O'Grady, M. (2018, August 10). Walmart Is Currently The World's Largest Retailer, But Alibaba And Amazon Are Closing In. Retrieved from https://go.forrester.com/blogs/walmart-is-currently-the-worlds-largest-retailer-but-alibaba-and-amazon-are-closing-in/

O'Kane, S. (2018, May 07). Uber reportedly thinks its self-driving car killed someone because it 'decided' not to swerve. Retrieved from https://www.theverge.com/2018/5/7/17327682/uber-self-driving-car-decision-kill-swerve

O'Neil, C. (2017). The era of blind faith in big data must end. Retrieved from https://www.ted.com/talks/cathy_o_neil_the_era_of_blind_faith_in_big_data_must_end

Reich, R. (2007). “Politics diverted” (pp. 168-208). In R. Reich, Supercapitalism: The transformation of business, democracy, and everyday life. New York: Knopf

Rubin, J. and Carmichael, B. (2018). Reset: Business and Society in the New Social Landscape. Series: Columbia Business School Publishing; Published by: Columbia University Press

Rudin, C., & Wagstaff, K. L. (2014). Machine learning for science and society. Machine Learning, 95(1), 1-9. doi:10.1007/s10994-013-5425-9

SciShow. (2015, February 19). *The Science of Anti-Vaccination*. [Video file]. Retrieved December 8, 2018, from https://youtu.be/Rzxr9FeZf1

Shulevitz, J. (2018, November 07). Alexa, Should We Trust You? Retrieved from https://www.theatlantic.com/magazine/archive/2018/11/alexa-how-will-you-change-us/570844/

The Human Toll of Instant Delivery [Audio blog interview]. (2018, November 26). The New York Times Podcast, The Daily. Retrieved from https://www.nytimes.com/2018/11/26/podcasts/the-daily/warehouse-workers-instant-delivery.html?action=click&module=audio-series-bar®ion=header&pgtype=Article

U.S. Census Bureau. (2016, March). Educational Attainment in the United States: 2015 (Report No. P20-578). Washington, DC: U.S. Government Printing Office. Retrieved from https://www.census.gov/content/dam/Census/library/publications/2016/demo/p20-578.pdf

Vanian, J. (2018, June 25). Unmasking A.I.'s Bias Problem. Retrieved December 8, 2018, from http://fortune.com/longform/ai-bias-problem/

Vogel, D. (2005). The market for virtue: The Potential and Limits of Corporate Social Responsibility. Brookings Institution Press: Washington, DC.

What Are The New Jobs In A Human Machine World? (2018, October 16). Retrieved from https://www.forbes.com/sites/insights-intelai/2018/07/17/what-are-the-new-jobs-in-a-human--machine-world/#3b1ce9d563e3

White House Hosts Summit on Artificial Intelligence for American Industry. (2018, May 10). Retrieved from https://www.whitehouse.gov/articles/white-house-hosts-summit-artificial-intelligence-american-industry/

Appendix A

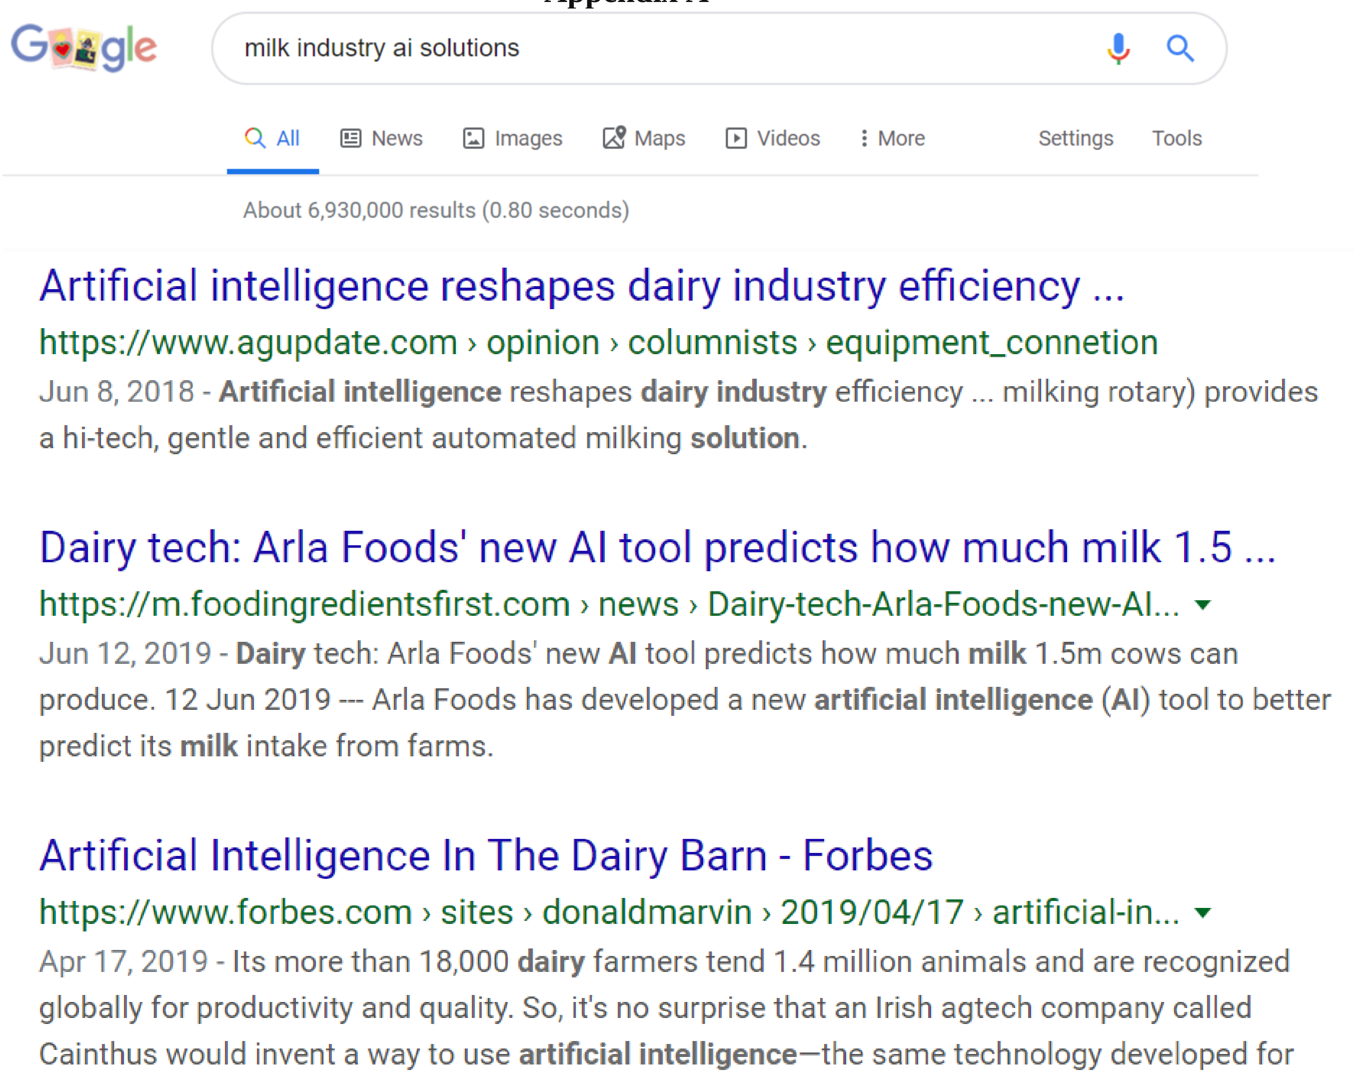
milk industry ai solutions

All News Images Maps Videos More Settings Tools

About 6,930,000 results (0.80 seconds)

Artificial intelligence reshapes dairy industry efficiency ...
https://www.agupdate.com › opinion › columnists › equipment_connetion
Jun 8, 2018 - **Artificial intelligence** reshapes **dairy industry** efficiency ... milking rotary) provides a hi-tech, gentle and efficient automated milking **solution**.

Dairy tech: Arla Foods' new AI tool predicts how much milk 1.5 ...
https://m.foodingredientsfirst.com › news › Dairy-tech-Arla-Foods-new-AI...
Jun 12, 2019 - **Dairy** tech: Arla Foods' new **AI** tool predicts how much **milk** 1.5m cows can produce. 12 Jun 2019 --- Arla Foods has developed a new **artificial intelligence** (**AI**) tool to better predict its **milk** intake from farms.

Artificial Intelligence In The Dairy Barn - Forbes
https://www.forbes.com › sites › donaldmarvin › 2019/04/17 › artificial-in...
Apr 17, 2019 - Its more than 18,000 **dairy** farmers tend 1.4 million animals and are recognized globally for productivity and quality. So, it's no surprise that an Irish agtech company called Cainthus would invent a way to use **artificial intelligence**—the same technology developed for

Appendix B

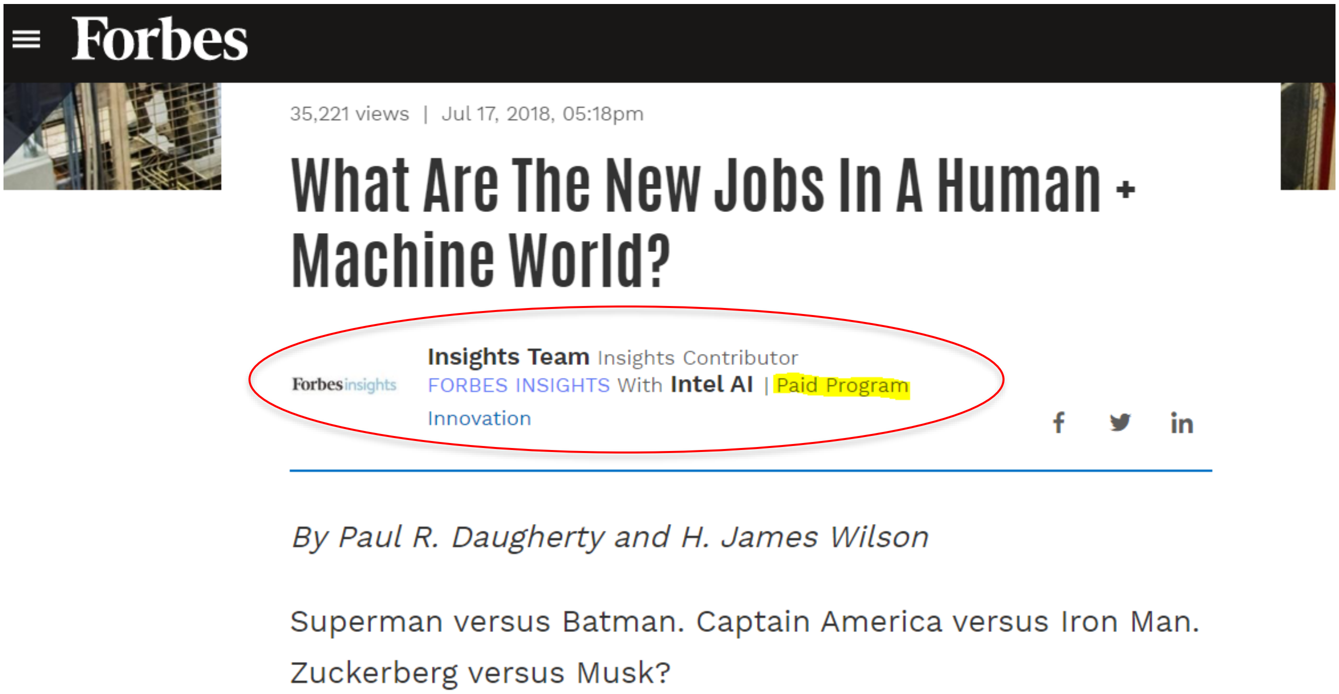

Forbes

35,221 views | Jul 17, 2018, 05:18pm

What Are The New Jobs In A Human + Machine World?

Forbes insights

Insights Team Insights Contributor

FORBES INSIGHTS With **Intel AI** | Paid Program

Innovation

By Paul R. Daugherty and H. James Wilson

Superman versus Batman. Captain America versus Iron Man. Zuckerberg versus Musk?

Appendix C

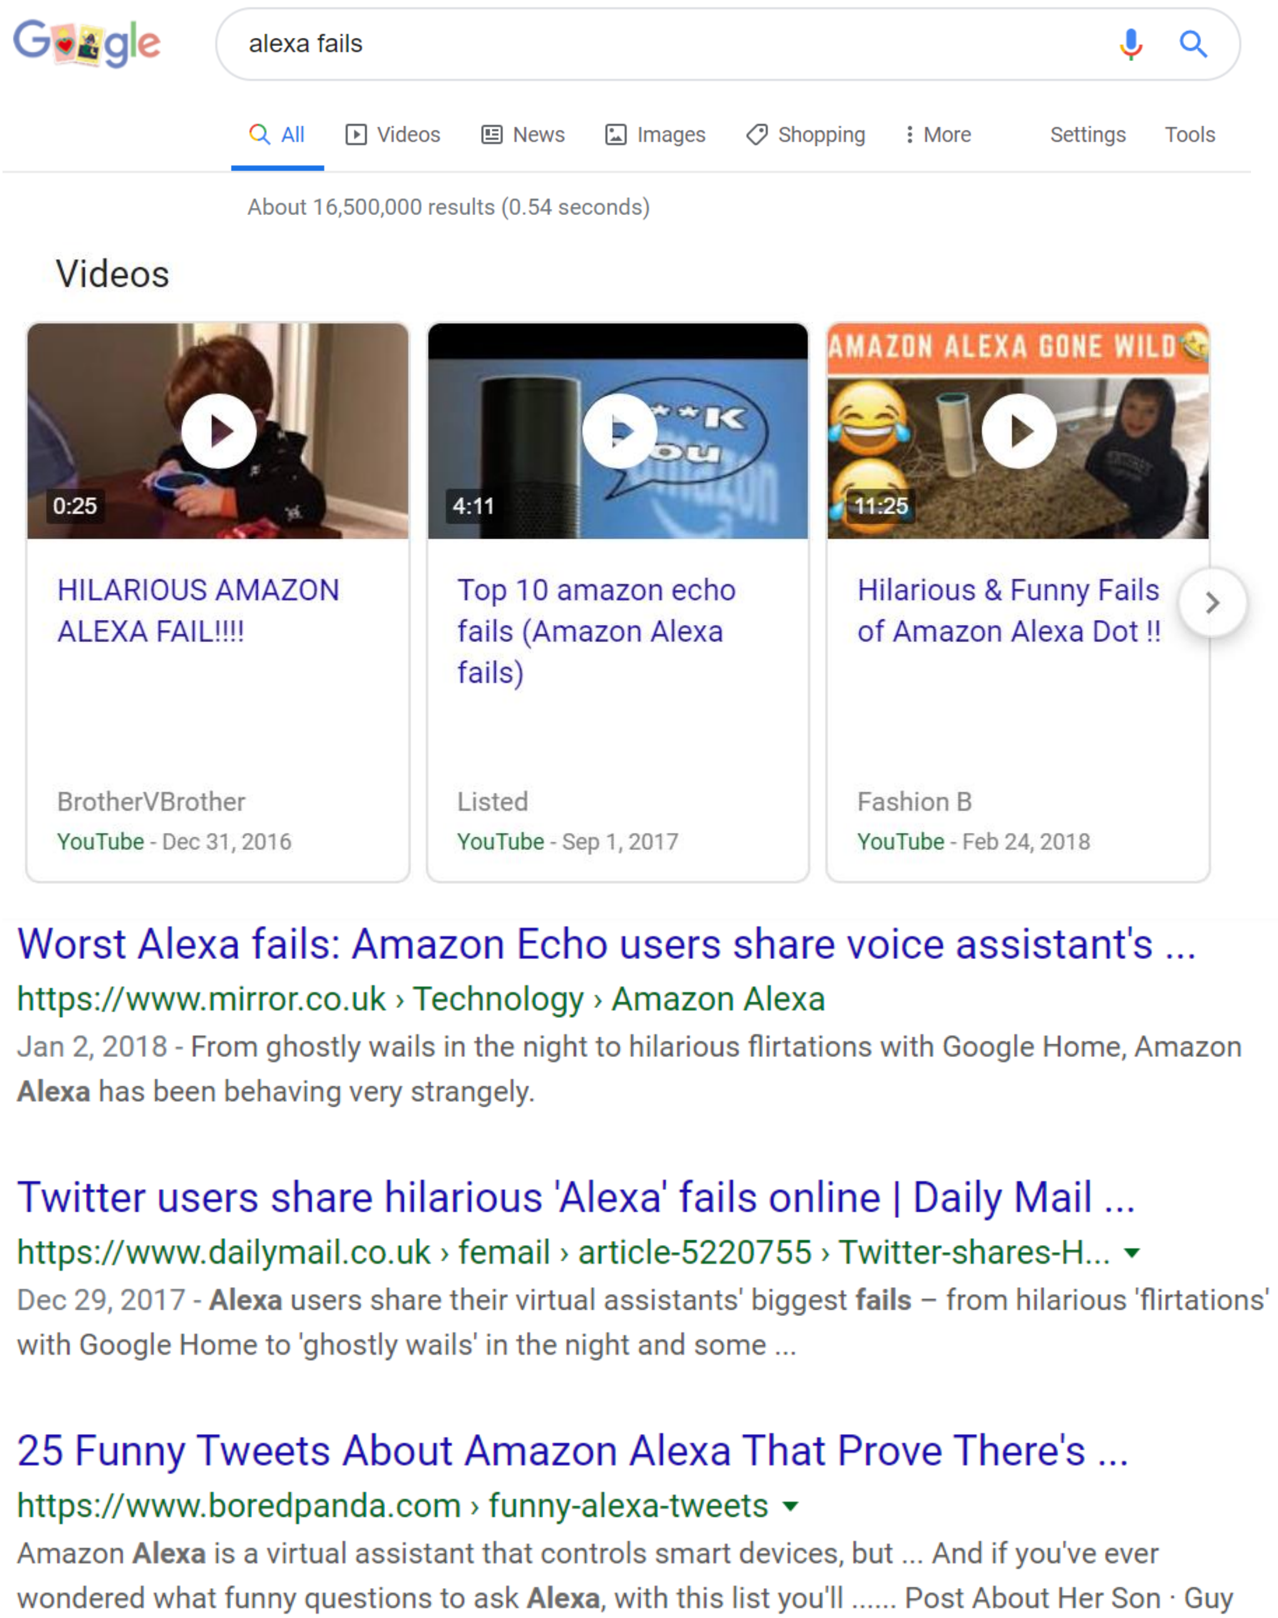
alexa fails

All Videos News Images Shopping More Settings Tools

About 16,500,000 results (0.54 seconds)

Videos

0:25 HILARIOUS AMAZON ALEXA FAIL!!!! BrotherVBrother YouTube - Dec 31, 2016

4:11 Top 10 amazon echo fails (Amazon Alexa fails) Listed YouTube - Sep 1, 2017

AMAZON ALEXA GONE WILD 11:25 Hilarious & Funny Fails of Amazon Alexa Dot !! Fashion B YouTube - Feb 24, 2018

Worst Alexa fails: Amazon Echo users share voice assistant's ...
https://www.mirror.co.uk › Technology › Amazon Alexa
Jan 2, 2018 - From ghostly wails in the night to hilarious flirtations with Google Home, Amazon **Alexa** has been behaving very strangely.

Twitter users share hilarious 'Alexa' fails online | Daily Mail ...
https://www.dailymail.co.uk › femail › article-5220755 › Twitter-shares-H...
Dec 29, 2017 - **Alexa** users share their virtual assistants' biggest **fails** – from hilarious 'flirtations' with Google Home to 'ghostly wails' in the night and some ...

25 Funny Tweets About Amazon Alexa That Prove There's ...
https://www.boredpanda.com › funny-alexa-tweets
Amazon **Alexa** is a virtual assistant that controls smart devices, but ... And if you've ever wondered what funny questions to ask **Alexa**, with this list you'll ...... Post About Her Son · Guy