GoSpotCheck

Ethical AI in the Enterprise

How fairness, transparency, and social science will predict the success of AI initiatives

Concept in Brief

The Problem:

AI promises to propel businesses forward, but questions of ethical practices and responsible design are holding them back from achieving lasting value.

The Reason:

Inherent human bias, discrimination, and lack of diverse data sets are misguiding even the most well-intentioned AI initiatives.

The Answer:

Adopting ethical design principles with human-centric strategies, assembling multidisciplinary teams, and creating strong forms of governance will ensure companies realize the full value of AI, and avoid financial and reputational risk.

AI: Demystified

For years, the progression of Artificial Intelligence has excited and intimidated many of the world's leading organizations. Named the single most disruptive technology by 69.5% of senior executives in 2020, AI and its subsets (such as Machine Learning, Robotic Process Automation, and Natural Language Processing) will continue to transform the way in which humans live, work, play, shop – and interact with other humans. Researchers believe the revolutionary outcomes of AI will be especially valuable to business, particularly when applied to operational efficiencies and automated business processes. As solutions increase in sophistication, pioneering companies will move swiftly toward enterprise-wide investment and adoption, as predicted by IDC and Forrester:

In 2020, 25% of the Fortune 500 will add AI building blocks to their Robotic Process Automation (RPA) efforts to create hundreds of new Intelligent Process Automation (IPA) use cases (e.g. text analytics and machine learning).[1]

By 2022, 75% of brands will embed intelligent automation into technology and process development, using AI-based software to discover operational and experiential insights to guide innovation.

In 2023, global spending on AI systems will hit $97.9B which is up from $37.5 billion in 2019. [2]

By 2024, 25% of spend will be as "Outcomes-as-a-service" with AI integral to every part of the business.

"One thing is for certain:

Not all AI is created equal, and not all AI relies on the same core technologies."

– Jeremy Fain, CEO and Co-Founder of the deep learning platform Cognitiv

Despite its anticipated growth, AI's seemingly limitless impact on organizations causes mixed feelings for even the most seasoned industry leaders – often stemming from a lack of understanding of what AI truly is. Jeremy Fain, CEO and Co-Founder of the deep learning platform Cognitiv, addresses five of the most common misconceptions associated with AI – the primary being the use of one universal definition. "Because the term 'AI' has grown to encompass so many different things, it's important to be clear about what type of AI one is referring to," says Fain. "One thing is for certain: Not all AI is created equal, and not all AI relies on the same core technologies."[3]

What is AI?

While organizations and luminaries may each have varying definitions of AI depending on their use cases and technologies, AI abstractly refers to the functional capability of machines to mimic human behaviors and human intelligence. Some of the most common subsets of AI include:

Machine Learning (ML)

Machine learning is an approach to data analysis that involves building and adapting models, which allow programs to "learn" through experience. Machine learning involves the construction of algorithms that adapt their models to improve their ability to make predictions.[4]

Where you've seen it: Uber, Netflix, Siri, Cortana, and Amazon Alexa

Robotic Process Automation (RPA)

Robotic process automation is a productivity tool that allows a user to configure one or more scripts (which some vendors refer to as "bots") to activate specific keystrokes in an automated fashion. The result is that the bots can be used to mimic or emulate selected tasks (transaction steps) within an overall business or IT process.[5]

Where you've seen it: Help desk software, chatbots, credit card application

Natural Language Processing (NLP)

Natural language refers to language that is spoken and written by people, and natural language processing (NLP) attempts to extract information from the spoken and written word using algorithms.[6]

Where you've seen it: Voice text messaging, language translation, autocorrect

These ground-breaking technologies may seem more machine than human – or, even more commonly misconstrued: entirely machine, without human input or ongoing supervision. However, a cornerstone of AI adoption in the enterprise is quite the opposite: the need for human training from initial deployment to ongoing functioning. Harvard Business Review's 2018 study, "Collaborative Intelligence: Humans and AI are Joining Forces," explored the dynamic between humans and machines with the rise of AI. In its study, HBR analyzed over 1,000 companies in 12 industries and identified five collaborative principles key to attaining business impact[7]:

1. **Reimagine** business processes
2. **Embrace** experimentation/employee involvement; actively direct AI strategy
3. **Responsibly** collect data
4. **Redesign** work to incorporate AI
5. **Cultivate** related employee skills.

Fig. 1: Examining the value of collaborative intelligence using Harvard Business Review's identified principles

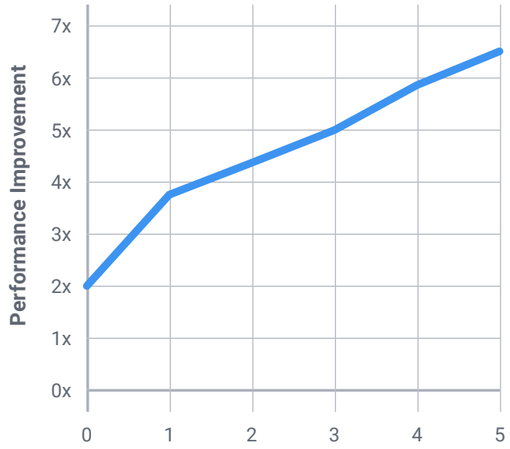

0 Indicates the adoption of only basic, noncollaborative AI

"The more of these principles companies adopted, the better their AI initiatives performed in terms of speed, cost savings, revenues, or other operational measures."

– S H. James Wilson and Paul R. Daugherty

When humans and machines work in tandem, their symbiotic relationship positions a business for long-term growth. Spurred by the potential of this collaborative intelligence, smart businesses are thoughtfully considering the traits and roles of humans assigned to training AI systems: Microsoft's Cortana, its AI assistant built to support the company's suite of products, required dedicated hours of training from a human team – which included a poet, a novelist, and a playwright – to develop its "confident, caring, and helpful but not bossy" demeanor to interact, please, and help human users. As a result, Cortana is the world's second-largest virtual assistant, with 22% in global market share.[8]

Similarly, Amazon Alexa and Apple's Siri have become icons of their era, even seeming eerily human at times; but they serve as reminders that behind each response is a team of human trainers, explainers, and sustainers – or workers constantly ensuring the AI is functioning properly and responsibly. Designing human-centric AI begins with a human-centric leadership philosophy, making hiring, recruiting, and people operations crucial to the deployment of advanced technology.

Instead of AI replacing human tasks – another common misconception among skeptics – it should complement and enhance existing human capabilities, while also making room for new roles. Without human input, a machine is merely bells and whistles: "As advanced as technology gets, it will always be learning from us rather than the other way around," Explains Adam Waytz, psychologist and associate professor of management and organizations at the Kellogg School of Management at Northwestern University.[10] As humans and AI collaborate to create value, knowledgeable assigned teams are needed to safely deploy and manage the system – so the future of AI lies in human hands, and is vulnerable to human shortcomings.

The Humans Behind the Machines: 3 New Developing Roles Driving AI Adoption

Trainers: Instructors who coach the AI on how to perform human tasks and exhibit human traits

Explainers: Translators who assist organizational leaders, colleagues, and customers in understanding the inner-workings of the AI

Sustainers: Guards who ensure the AI is functioning within governmental and organizational guidelines[9]

The Inevitability of Ethics in AI Adoption

Humans are behind the wheel of AI implementation. As AI evolves, so will the risks of human misuse, bias, and discrimination – even with the most well-intentioned projects. Forrester identified 2020 as the year in which implementation and exploration of AI will accelerate significantly; but the research firm also predicted consequential setbacks associated with adoption, including multiple high-profile PR disasters. "AI can perpetuate harmful discrimination, bias, and even lead to fatal consequences. The spread of deepfakes, misuse of facial recognition, and overuse of personalization can harm, offend, or creep out customers and employees."[11] "The result is that AI will benefit some of us far more than others, depending upon who we are, our gender and ethnic identities, how much income or power we have, where we are in the world, and what we want to do." says S.A. Applin, Ph.D.

By 2023, Gartner predicts one-third of all brand public relations disasters will result from data ethics failures.[12]

Despite the popularity of voice assistants like Cortana, Siri, and Amazon Alexa, questions of gender bias are increasing as a result of the products' docile female voices – a conscious choice made by each company to appeal to psychological human preference. And although AI can technically be genderless, each company settled on a "helpful, supportive, and trustworthy" female tone, which proved to be the preferred choice when tested among humans before launch. In contrast, IBM's Watson, a cognitive AI assisting humans in technical sectors like healthcare and law, and even claiming victory in Jeopardy, speaks in a male authoritative tone. This product design choice was meant to communicate a high standard of expertise compared to daily voice assistants.[13] Nevertheless, with more than 90 million smartphone users tapping into voice-enabled AI at least once a month,[14] the pervasive role of these virtual assistants is surfacing sociological debate around humans' biased decision-making. Perhaps even more concerning to the public are the reports of mass "listening" by devices in the background of users' lives. In 2014, the story of Ruthy Hope Slatis and her new job transcribing audio files for Amazon's "top-secret speech-recognition product" made headlines. The product was revealed to be none other than the Echo featuring Alexa, and the audio files Slatis used to power the project had been snippets of random, intimate moments from inside the homes of unsuspecting people.[15] Soon after, Slatis quit, and though Amazon announced its speakers only record audio once launched by the user, stories like these naturally cause unease around human rights of privacy and trust. In the same vein, over-personalization and breaching of privacy as a result of companies' marketing efforts will only continue to turn away consumers: "73% of consumers in the United States, Canada, Japan, Australia, France and the United Kingdom think people using connected devices should worry about eavesdropping, and 63% think connected devices are 'creepy' in the way they collect data about people and their behaviors."[16] And one of AI's most infamous inventions, facial recognition – a technology originally designed to increase security in public spaces and even to enhance QA in food production facilities (where missing hats and hair-nets can be picked up by cameras) has produced racial bias. In a recent study performed by M.I.T. Media Lab, facial recognition misidentified 35% of images containing dark-skinned women, while calculating 99% correctly for white-skinned males.[17] The reason, according to experts, is dependent on humans' non-diverse sets of training data used to inform the technology.

"Training AI models to emulate human preferences or biased decision-making will only replicate the shortcomings of our own mind."

– Tomas Chamorro-Premuzic, Professor of Business Psychology at University College London (UCL) and Columbia University[19]

Even seemingly harmless Business Process Management solutions designed to increase efficiency in recruiting, hiring, and numerous workflow automations have stirred controversy associated with cyber-snooping and unintentionally-biased predictions.[18]

Although these examples represent only a fraction of the challenges facing the pioneers of AI today, they surface the questions we will grapple with over the coming decades as the technology proliferates. Humans are fundamentally required to implement and manage AI; but individual human flaws, biases, and demographic differences will inherently impact the design and performance of AI systems.

Common Controversial Forms of AI

Deepfakes: Deepfakes refer to manipulated videos, or other digital representations produced by sophisticated artificial intelligence, that yield fabricated images and sounds that appear to be real.

Facial Recognition: A facial recognition system is used to identify or verify a person from a digital image or a video frame from a video source, commonly used by law enforcement to enhance public safety and security.

Personalization: Personalization tools powered by AI allow brands to reach consumers with highly-relevant offers based off of their activity, demographics, and shopping patterns.

It's All in the Approach

Despite its risks, businesses should not be deterred from implementing AI – instead, AI will force organizations to think more holistically about not only their deployment strategies, but also about who and how they enroll humans in the future. According to Harvard Business Review, companies who wait for AI to mature and for technological expertise to grow will likely never catch up, due to AI's unique and time-consuming implementation requirements.[20] And for those early adopters, responsible data collection, ethical training, and the creation of knowledgeable teams will ultimately determine their success. "Whether for companies or countries, the greatest risk around AI today is one of exclusion," says Amir Husain, Founder & CEO of the global AI company, SparkCognition, and CEO of SkyGrid. "Those who fail to understand and embrace AI leave themselves vulnerable to an insurmountable disruption by those who come to grips with it first." Leading adopters should act quickly – but shouldn't sacrifice a wise, thoughtful, and ethical approach in the name of progress. If they do, they'll risk organizational repercussions that will impact performance in the long-term. The best and brightest will demonstrate an approach that's swift but thoughtful, outcome-oriented but human-centric. The challenges of AI are the most complex of any technology to come before it. And while there isn't one simple answer to solve the ethical challenges AI presents, some of the world's most renowned companies are beginning to create highly-responsible systems – and today's adopters should follow suit.

Using AI for Good

With such high stakes, organizations benefit from external experts who can lead a process to establish internal guidelines for ethical AI strategy and comport. In 2018, Microsoft established operational AI principles to ensure responsible data collection and management, identifying six fundamental requirements in system development that other corporations can replicate: [21]

Fairness
AI systems should treat all people fairly

Inclusiveness
AI systems should empower everyone and engage people

Reliability & Safety
AI systems should perform reliably and safely

Transparency
AI systems should be understandable

Privacy & Security
AI systems should be secure and respect privacy

Accountability
AI systems should have algorithmic accountability

To ensure new products adhere to these principles, Microsoft appointed the AETHER Committee (AI and Ethics in Engineering and Research), whose members include top research, engineering, ethics, law, and policy makers. The committee is charged with making formal recommendations on policies, processes, and best practices.[22] Other tech leaders are taking a similar approach, like Apple, Amazon, Sony, and Google, who have joined the Partnership on AI: a global, multi-stakeholder organization spanning 13 countries and four continents to educate and cultivate awareness of AI's effects on businesses and the general public. "We believe that this work must be informed by diverse voices, multidisciplinary expertise, and a deep understanding of the context in which technology is being used,"[23] says Terah Lyons, Founding Executive director of PAI.

In addition to these specialized committees of experts, the need for expansive AI knowledge inside and outside of enterprise is growing rapidly. Still nascent, many companies are exploring what kind of expertise will best suit these new roles. In heartening news for Gen Z, or True Gen members, the role of Artificial Intelligence Specialist (inclusive of developers and machine learning engineers) was recently named number one in LinkedIn's 2020 Emerging Jobs Report, with 74% annual growth since 2016. In a field poised for exponential growth, experts are advising organizations to consider the balance of education and human expertise needed, in addition to technical proficiency, to truly reap the benefits of this emerging multidisciplinary role.

In 2017, Barbara J Grosz, computer scientist and Higgins Professor of Natural Sciences at Harvard University, introduced the university's first ethical computer science course to "instruct the people who will build future AI systems in how to identify and think through ethical questions."[24] The course, "Intelligent Systems Design and Ethical Challenges," combines Grosz's knowledge of linguistics, philosophy, psychology, economics, anthropology, and sociology with the instruction of AI design. Now, Harvard University offers 18 different computer science courses with an ethical concentration, ranging from systems programming, machine learning and its effects on fairness and privacy, robots and work, to social networks and the question of censorship, and human-computer interaction. Likewise, the introduction of FAT ML, or Fairness, Accountability, and Transparency in Machine Learning has earned its own dedicated area of study within computer science, better-equipping tomorrow's AI developers and ML engineers with the diverse skill-set needed to design responsible solutions.

"This is why tech companies' AI labs need social science and cross-cultural research: It takes time and training to understand the social and cultural complexities that are arising in tandem with the technological problems they seek to solve." – S.A. Applin

Just as machines require humans to operate effectively, humans require a deep understanding of the sociological implications of AI in addition to technological expertise. The future of successful AI at enterprise scale relies heavily on preliminary understanding and adoption of responsible practices, many of which organizations can begin to incorporate today.

Adopting Ethical AI in the Enterprise

The first step may be the hardest: understanding the importance of ethical AI in your organization – and grasping the consequences of irresponsible systems if data is mismanaged. By accepting the complexities and challenges that come with ethical AI adoption, you've already made progress toward designing a human-centric strategy that may take years to develop. With the right tools, the right teams, and the right mindset, you can look to the future with confidence while paving the way for sustainable, scalable, and responsible growth. Here are the steps experts recommend to scale AI in the enterprise.

Build a Dedicated Team

Consider your current AI teams and their knowledge of ethics in AI systems: Are your engineers, specialists, and developers also experts in fairness and moral design? Developing highly-trained teams with experience in responsible AI – from leadership to engineers – is key to delivering unbiased results. Examples of roles emerging in today's enterprise worth considering include:

- Chief AI Ethics Officer
- AI Ethics Advisory Board
- Trainers, Explainers, and Sustainers with Ethical Computer Science/FAT ML expertise

Specify Goals & Start Small

While exploring AI technology partners and developing strategic initiatives, define tests concretely and confirm the ROI sought: increased efficiency, reduced risk, or increased revenue are all great goals. It's critical to understand how AI will be measured and scaled in a "test and learn" approach, and for there to be institutional frameworks in place to deploy it at scale once proven. Also key to this is executive alignment and de-siloing of data and analytics. Once your goals are established, start small with a test dataset, and remain resilient when failures arise. Testing and learning is all part of the process, but ensuring the system is functioning properly, safely, and ethically is critical before implementation on a larger scale.

Establish Guidelines & Governance

There is no universal code for ethical AI, which is where bias and demographic influence become obstacles. Establishing a set of regulations for the use of AI – much like Microsoft's six guiding principles – will give leaders a practical framework for decision-making. In addition, organizations should establish necessary degrees of oversight to ensure proper governance of programs, internally and externally. Transparency is critical.

Establish Transparency

Once AI is functioning in the enterprise, measuring results is the next hurdle. Data scientists are critical to AI initiatives and should be empowered to structure tests, combine data sets, measure results, and translate their findings to enable rapid decision making. Mechanisms are needed to initiate ongoing tests, democratize insights, and optimize systems to ensure growth across all lines of business.

Promote Continued Education

The study of responsible AI is still in a nascent stage. Organizations must encourage the ongoing study of their AI progression and program successes, and provide resources for company leaders, stakeholders, and relevant teams to be knowledgeable and prepared for change. Use PAI's Closing Gap Ideation Game as a resource to identify pitfalls and design solutions throughout implementation.

Artificial Intelligence will transform the lives of those who use it, train it, and those who benefit from it – from automating and simplifying business processes, to creating closer relationships with consumers.

These powerful capabilities mean even greater responsibility for human adopters, from initial system training to ongoing maintenance and measuring. The collaborative intelligence of humans and machines will continue to pose challenges for adopters in avoiding bias, promoting ethical responsibility, and understanding the technology's effects on the greater good. Those who approach AI implementation swiftly but wisely, and who place humanity at the center of their strategy will position themselves for a profitable future – one that's more human than we may think.

GoSpotCheck

GoSpotCheck is reimagining how tomorrow's workforce works, with software to power the people. The world's top brands distribute tasks to their frontline on our mobile app, then spot and close execution gaps using our business intelligence dashboards.

Our cloud platform presents the right action at the right time to the right person to improve execution and helps leaders make better decisions, faster.

- Direct tasks and execution on mobile
- Capture data efficiently with image recognition, powered by proprietary ML
- Ensure you're delivering what customers want–every moment, in every location
- Optimize revenue, labor, margin, assortment, and merchandising
- Reduce risk and QA breaches
- Have data-driven conversations with frontline teams and partners

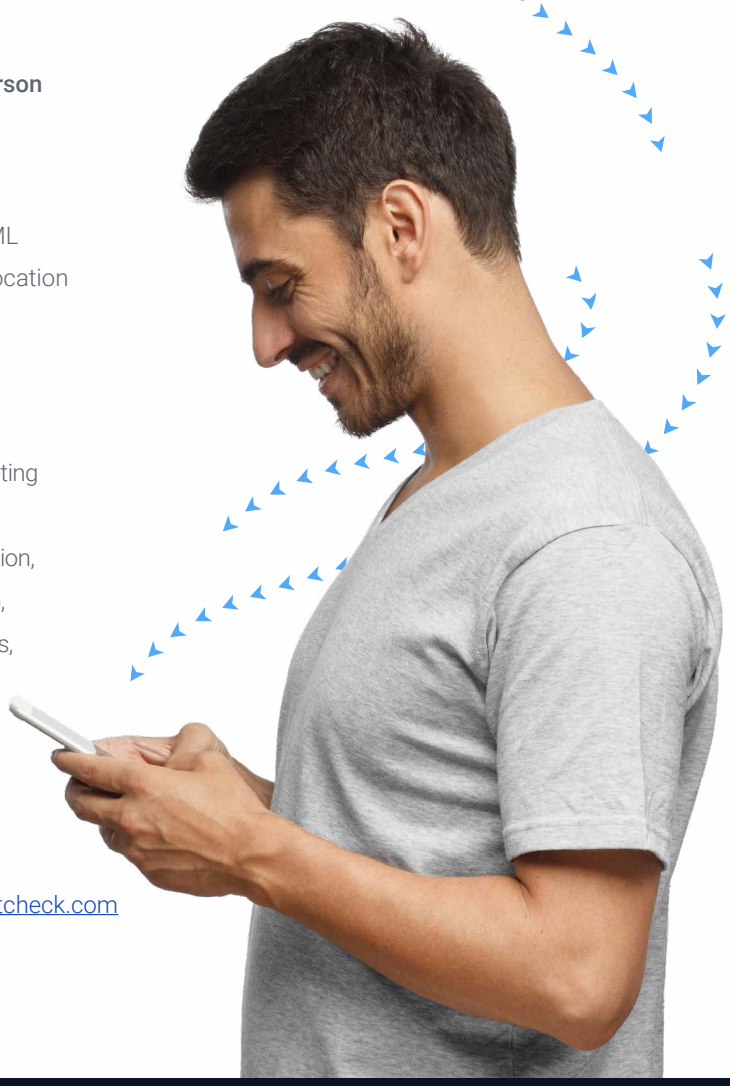

Enterprise brands use GoSpotCheck in 70 countries to manage field sales, marketing and merchandising, QA, asset management, competitive analysis, and training. We enable execution at scale with world-class task management, image recognition, and machine learning capabilities, for category leaders including ABInBev, Diageo, Beam, PepsiCo, Frito-Lay, Danone, Dole, Panera, Red Robin, Coffee & Bagel Brands, Dairy Queen, Save A Lot, Best Buy, Hudson News, and Levi's.

We enable complex business processes, with a talented team of 135 and live customer support, 7 days a week, from Denver, CO. Pricing is available in per license, per location, or enterprise packages.

For more information: gospotcheck.com, 844-359-2502 or solutions@gospotcheck.com

Learn More

References

[1] https://www.forbes.com/sites/gilpress/2019/11/22/top-artificial-intelligence-ai-predictions-for-2020-from-idc-and-forrester/#26936d8e315a

[2] https://www.zdnet.com/article/spending-on-ai-systems-set-to-surge-and-consultants-are-going-to-thrive/

[3] https://www.forbes.com/sites/forbesagencycouncil/2019/03/08/five-things-most-people-get-wrong-about-ai/#674ed9302ad8

[4] https://deepai.org/machine-learning-glossary-and-terms/machine-learning

[5] https://www.gartner.com/en/information-technology/glossary/robotic-process-automation-rpa

[6] https://pathmind.com/wiki/natural-language-processing-nlp

[7] https://hbr.org/2018/07/collaborative-intelligence-humans-and-ai-are-joining-forces

[8] https://futuresource-consulting.com/reports/?year=2019

[9] https://www.fastcompany.com/40563782/how-a-i-anxiety-is-creating-more-jobs-for-humans

[10] https://www.cnn.com/2019/05/01/perspectives/milken-conference-ai-workers-ethics/index.html

[11] https://go.forrester.com/blogs/predictions-2020-ai/

[12] https://www.gartner.com/en/newsroom/press-releases/2019-12-02-gartner-predicts-80--of-marketers-will-abandon-person

[13] https://www.pcmag.com/opinions/the-real-reason-voice-assistants-are-female-and-why-it-matters

[14] https://www.vox.com/future-perfect/2019/6/12/18660353/siri-alexa-sexism-voice-assistants-un-study

[15] https://www.bloomberg.com/news/features/2019-12-11/silicon-valley-got-millions-to-let-siri-and-alexa-listen-in

[16] https://www.forbes.com/sites/gilpress/2019/06/09/9-indicators-of-the-state-of-artificial-intelligence-ai-may-2019/#4867f9b5577f

[17] https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html

[18] https://www.forbes.com/sites/tomaspremuzic/2018/05/27/four-unethical-uses-of-ai-in-recruitment/#4e92da4015f5

[19] https://www.forbes.com/sites/tomaspremuzic/2018/05/27/four-unethical-uses-of-ai-in-recruitment/#4e92da4015f5

[20] https://www.forbes.com/sites/amirhusain/2019/06/26/the-only-thing-we-have-to-fear-is-the-fear-of-ai/#1bddde8e2593

[21] https://www.microsoft.com/en-us/ai/responsible-ai

[22] https://www.microsoft.com/en-us/ai/responsible-ai

[23] https://www.partnershiponai.org/wp-content/uploads/2021/01/PAI-2019-Annual-Report.pdf

[24] https://harvardmagazine.com/2019/01/artificial-intelligence-limitations