ACCENTURE RESPONSE TO THE FEDERAL TRADE COMMISSION

"COMPETITION AND CONSUMER PROTECTION ISSUES ASSOCIATED WITH THE USE OF ALGORITHMS, ARTIFICIAL INTELLIGENCE, AND PREDICTIVE ANALYTICS IN BUSINESS DECISIONS AND CONDUCT" HEARINGS INITIATIVE

FEBRUARY 6, 2019

Background on Algorithms, Artificial Intelligence, and Predictive Analytics, and Applications of the Technologies

....

What features distinguish products or services that use algorithms, artificial intelligence, or predictive analytics? In which industries or business sectors are they most prevalent?

....

AI is what economists call a *general-purpose technology*—it is a collection of different technologies that allow smart machines to extend human capabilities and intelligence by sensing, comprehending, acting and learning—thereby allowing people to achieve much more than can be achieved without the technology.

Like electricity and the internal combustion engine, AI's significance lies in the fact that it will cause significant transformations not only through direct contributions to society, but also through the way its spill-over effects enable a vast range of complementary innovations. Electricity made possible factory electrification, telegraphic communication, and all that followed. And the internal combustion engine gave rise to the automobile, the airplane, and modern transportation and logistics networks.

Our framework for understanding AI, and to get a sense of the technologies that underlie it, is centered around four principle qualities that enable a machine to:

- *Sense.* AI lets a machine perceive the world around it by acquiring and processing images, sounds, speech, text, and other data.
- *Comprehend*. AI enables a machine to understand the information it collects by recognizing patterns, much as humans interpret information by understanding the patterns presented and their context. However, a machine does not derive true "meaning".
- *Act.* AI can enable a machine to take actions in the physical or digital world based on that comprehension.
- *Learn.* AI enables a machine to continuously optimize its performance by learning from the success or failure of those actions.[i]

To understand how AI features into analytics we must think of analytics as the field of analyzing data to improve decision making. The goal of analytics is to derive insights from data. Analytics, and the AI algorithms that support it, can have different levels of sophistication depending on the degree of insight required:

- At the basic end of the scale are so-called "*descriptive analytics*" which analyze historical data to understand what happened and why.
- More complex are "*predictive analytics*" which use data to predict what will happen in the future.
- Most complex are "*prescriptive analytics*" which not only forecast what will happen but tell you what you need to do about it.[ii]

We have recently moved beyond research and development to the real-world application of AI. This shift has been driven by the combination of greater and more affordable computing power and storage, the growth in data volumes, and the rise of open source frameworks. Early-stage consumer AI applications, which have one or more of the principle qualities above, are an increasingly normal part of everyday life, with tools such as Apple's Siri and Google Now on our mobile devices, Amazon's Alexa in our home, and chatbots online now part of our everyday customer experience. We have, however,

become much more recently aware of 'AI' in the past year, thanks, in no small part, to the significant increase in media attention.

These four categories are the ways in which AI can drive growth and improve the way we work:

1. *Furthering innovation*- Artificial Intelligence can stimulate additional innovation and technological change. It will require new organizational practices and institutions that increase the efficiency of production.
2. *Intelligent automation*- AI-powered innovation automates complex tasks such as analyzing large amounts of data and tasks much better suited for a machine: routine tasks and repetitive tasks. This frees up the worker to provide more value add to the company through the capabilities that are uniquely human: creativity, judgement, and improvisation.
3. *Increased contribution of capital*- AI can make the current capital stock, such as machines, more productive. If machines are shown to be more productive, this can lead to increased investment in capital thereby increasing the overall capital stock.
4. *Augmented Labor*- AI can amplify and extend human skills. Human-Machine collaboration will result in dramatic gains in business productivity. With the help of AI, people will have access to a wider range of insights, allowing them to make more informed decisions at a much larger scale.[iii]

Over the past twenty years or so, business has faced its fair share of disruption. Indeed, digital disruption is thought to have put more than half of Fortune 500 companies out of business since 2000. And now AI is set to compound that disruption by shifting it up to the next gear. AI will impact society on a similar scale. The immediate and most tangible outcome of these developments is the automation and augmentation of complex knowledge tasks within the enterprise, using AI software that can learn dynamically from data and interactions.

The impact of AI, once fully realized, could be significant. Accenture found that AI has the potential to double economic growth rates by 2035, changing the nature of work and boosting labor productivity by up to 40 percent. Indeed, far from creating a dehumanized future, our research affirmed that the real promise of AI is its power to augment what humans are great at and make them better at what they do. We're already seeing AI being embraced by companies in every industry. In Financial Services, for example, robo-personalized advisors, intelligent underwriting and automated agent-based banking services are already in action and will soon be commonplace. Life Sciences and health care companies are leveraging AI's power in areas including drug discovery and pharmacovigilance. And in the retail sector, companies now use intelligent recommendation systems and in-store experience enhancement.[iv] We modeled the economic impact of artificial intelligence (AI) in 16 industries, across 12 countries:
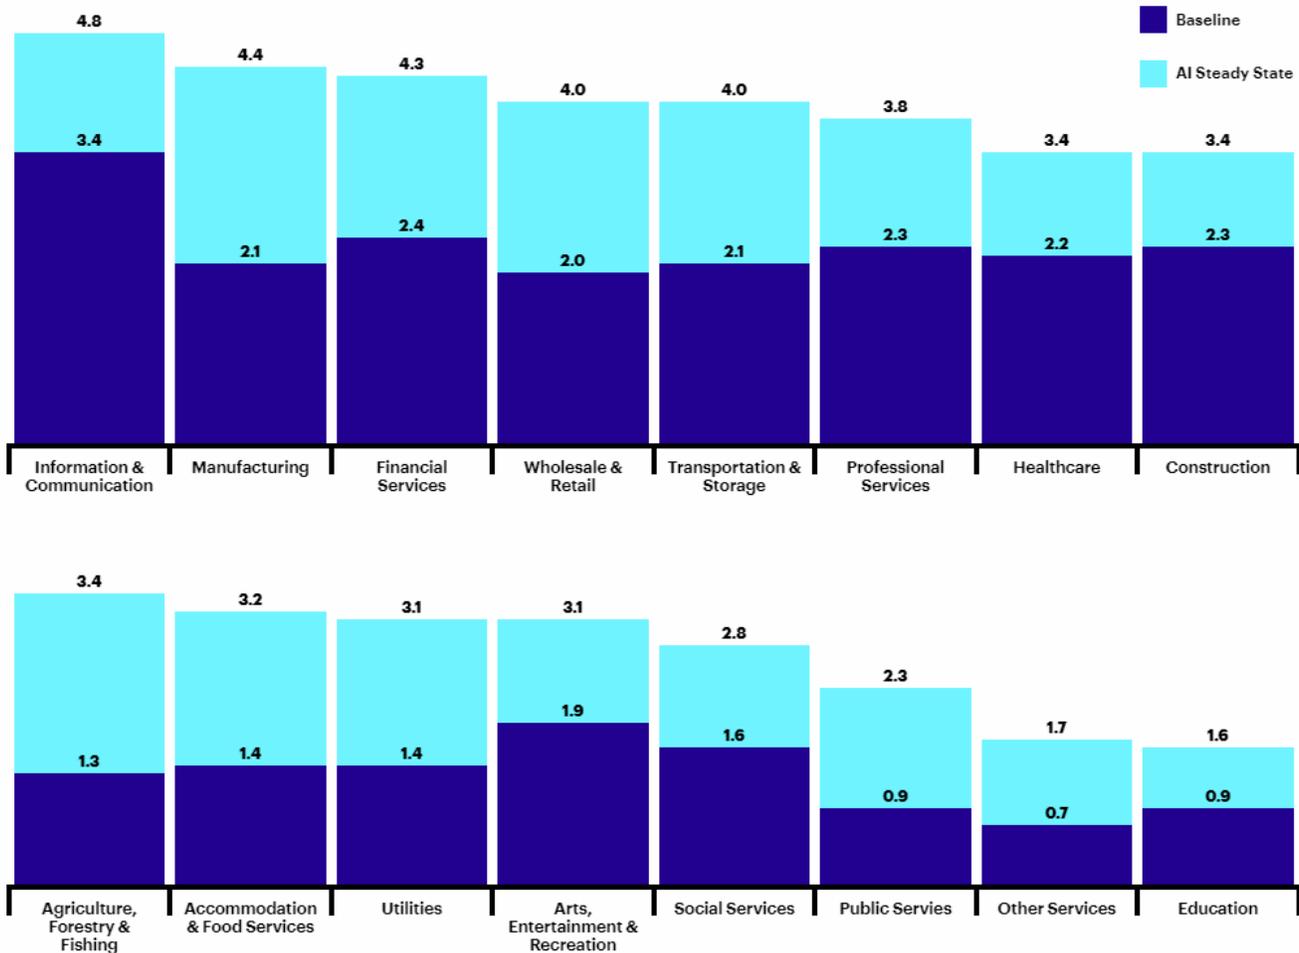
Annual growth rates in 2035 of gross value added (a close approximation of GDP), comparing baseline growth in 2035 to an artificial intelligence scenario where AI has been absorbed into the economy.
Source: Accenture and Frontier Economics

What factors have facilitated the development or advancement of these technologies? What types of resources were involved (e.g., human capital, financial, other)?

The progress of AI has accelerated in recent years due to the rise of big data, the proliferation of connected sensors and actuators (the Internet of Things), and access to vastly increased and cheap processing power, especially through the cloud. More than a productivity enhancer, we view AI as an entirely new factor of production: by optimizing processes with intelligent automation systems, by augmenting human labor and physical capital, and by propelling new innovations.

But to capitalize on this tremendous opportunity, businesses in every industry will need to synthesize AI into their strategies and develop responsible AI systems. Responsible AI systems are aligned to society's moral and ethical values and provide a clearer understanding regarding better outcomes as well as trade-offs. They take into consideration key issues such as: privacy, liability, transparency, bias, human values, data protection, and cybersecurity. Fundamentally, a company's AI deployments need to be aligned to a company's core values and ethical principles to benefit and engender trust with consumers, employees, business, and all of society.

AI's unique characteristics as a capital–labor hybrid—such as the ability to augment human labor at scale and speed, self-learn and continuously improve over time—will require new approaches and models in areas such as investment, innovation and human capital development.[v]

....

Are there factors that have impeded the development of these technologies? Are there factors that could impede further development of these technologies?

....

Companies are looking to AI to bring new skills and tools that help their people do new jobs and invent a new future for themselves and the organization. AI adopters are being confronted with some tough decisions in order to scale AI for the enterprise. These include:

- *Data residence and security* - Companies may not be comfortable sharing sensitive data in public or with third party administered clouds (especially if financial and healthcare-related personally identifiable information is involved). This poses challenges for machine learning and data analytics. This is one of many challenges related to the lack of a clear data privacy policy framework in the United States.
- *Customizability and ease of use* - Because service end-points can often be opaque and relatively inflexible, it can be challenging to customize AI solutions/modules to business needs.
- *Technology obsolescence* - AI is a rapidly maturing and fast changing field, in which new, disruptive technologies are emerging almost every week; some companies may feel that the risk of technological obsolescence runs higher when they are tied to a single vendor rather than having a rich ecosystem of providers.
- *Future proofing* - The latest AI technologies, such as deep learning, are growing in sophistication at a rapid pace: the evolution from conventional neural networks to word2Vec to deep recursive/recurrent/ convolutional networks took place within just over ten years. Companies need to be consistently ahead of the curve to exploit the potential of best-of-breed technologies, and that's not possible when they're tied to a single vendor that may become legacy at any moment.
- *Insufficient coordinated R&D investment* – In a recent submission to the National Science and Technology Council (NSTC) Select Committee on Artificial Intelligence, Accenture recommended that the federal government define and increase AI R&D investments; create a formal engagement strategy between the public and stakeholders to better identify and monitor investments; and anticipate risks and challenges for long-term R&D.[vi]

In many ways, these challenges flow from a marketplace dynamic that's currently out of step with the breakneck pace of technology change in AI, as well as the spend limits within which IT departments must operate today.

Challenges for companies also go beyond IT. A global survey of business leaders, commissioned by Accenture Applied Intelligence, Intel and SAS, conducted by Forbes Insights in September 2018, shows just how widely AI is being used by organizations globally. This survey reveals the significant momentum behind AI and leaves little doubt that it's poised to emerge as a transformative force in business, government and society. It also underlines the urgent need for guardrails to guide and manage its adoption, which, if not managed properly, could impede development and adoption.[vii] These challenges include:

- *Management and Oversight* - Despite widespread messages in popular media that suggest that AI operates independent of human intervention, those responsible for putting AI to work in their organizations across industries recognize that these technologies require rigorous management and oversight. In fact, nearly a quarter (23 percent) of AI adopters review or evaluate AI outputs at least daily. The data also suggest that companies that have been more successful with AI tend to have more rigorous management and oversight processes in place. For example, 74 percent of successful companies report that they review their AI output at least weekly, compared

with 33 percent of those that are less successful. And 43 percent of successful companies have a process in place for augmenting or overriding questionable results, compared with 28 percent of companies that haven't yet found success in their AI initiatives.[viii] Consideration should be given to global, industry driven guidelines, standards and best practices that can help create and safeguard trust at the heart of AI-driven systems and business models and permit the flexibility for innovation, allowing codes to develop with the technology. Examples include the IEEE Global Initiative for Ethical Considerations in AI and AS and its nine pipeline standards on Ethically Aligned Design.

- *Ethics-* In addition, AI adopters indicated relatively strong ethical processes in place today, with 63 percent affirming that they "have an ethics committee that reviews the use of AI," and 70 percent indicating they "conduct ethics training for their technologists." Among companies that report successful achievements from their deployment of AI, 92 percent say they conduct ethics training for their technologists, compared with 48 percent of those that have not achieved real success yet from deploying AI.[ix]
  However, there is uncertainty – 45% of respondents agree that not enough is understood about the unintended consequences of AI and 88% of respondents do not have confidence in AI-based decisions and output. In addition to management and oversight, organizations should consider building ethics into the innovation process. Data Science-based tools, methods, and capabilities can help address concerns around Responsible AI. Existing solutions are centred around the design-build-operation of AI systems, and includes the Accenture Fairness Tool, which is designed to understand and enable humans to better see and address bias in both the data and the algorithmic models that are at the core of AI systems.

- *Workforce -* Finally, perception of the impact of AI on our workforce represents a potential impediment to not just AI development and implementation, but also workforce development policies that can facilitate the creation and augmentation of jobs and the generation of new businesses. In addition to preparing and developing this future workforce, the U.S. faces current shortages in highly skilled professionals in science, technology and mathematics, among others, that could hold back the innovation of intelligent technologies.

Nearly 60 percent of respondents to our survey either "strongly" or "completely" agreed with the statement "we do not anticipate any impact on jobs due to AI's implementation." It comes as no surprise, then, that similar numbers indicated that they are "already seeing job roles being elevated as a result of AI (i.e. it will allow for employees to focus on more strategic rather than operative tasks) – 64 percent either "strongly" or "completely" agreed with this statement. They don't see jobs being replaced by AI – they see them being elevated. [x] At Accenture, for example, we have developed a technology platform to streamline and automate processes in areas such as finance and accounting, marketing, and procurement. For the hundreds of thousands of workers whose jobs have been impacted by this new platform, we provided internal resources and retrained them for other positions within the company. We also assist clients in undertaking similar transformations.

Wanted: Artificial intelligence experts

In artificial intelligence, job openings are rising faster than job seekers.

AI-related jobs include machine learning engineer, predictive modeler, corporate analytics manager, data scientist, computer vision engineer, computational linguist, and information strategy manager. | Source: Indeed.com | Credits: Ann Saphir, Data Visualization Engineer, Reuters[xi]

....

What are the advantages and disadvantages for consumers and for businesses of utilizing products or services facilitated by algorithms, artificial intelligence, or predictive analytics?

....

Accenture research shows AI holds vast potential for ushering in a new era of growth for businesses across industries. Accenture, in association with Frontier Economics, modeled the potential economic impact of AI for 16 industries covering a diverse field, from Manufacturing to Utilities to Healthcare. As our yardstick, we used growth in gross value added (GVA), a close approximation of GDP. GVA is an output measure that accounts for the value of goods and services produced in a certain sector. It can be thought of as the contribution of different sectors to economic growth. We compared two scenarios for each industry to assess AI's future impact.

First, the baseline case shows the expected economic growth for industries under current assumptions, which do not fully factor in AI development and implementation across the targeted industries and economic processes. Second, the AI steady state shows expected growth with AI integrated into economic processes. As it takes time for the impact of a new technology to feed through, we used 2035 as the year of comparison (for further details see "Appendix: Modeling the GVA impact of AI"). Our research shows that, across 12 countries, Information and Communication, Manufacturing and Financial Services are the three sectors that will see the highest annual GVA growth rates in an AI scenario, with 4.8 percent, 4.4 percent and 4.3 percent respectively by 2035 (Figure 4). In the Information and Communication industry, with its heavy reliance on technologies, AI capabilities can coalesce with existing systems to generate US$4.7 trillion in gross value added in 2035 (Figure 5). For instance, providers can develop new AI platforms for offering cyber-attack protection services to their customers.

In Manufacturing, precursors like the Internet of Things (IoT) create favorable conditions for the seamless integration of intelligent systems. Today's IoT technologies enable physical devices such as assembly lines to connect and communicate with digital systems. Moreover, AI can bridge the gap between current forms of automation and learning with more advanced forms. Our research shows that AI could add an additional US$3.8 trillion GVA in 2035 to this sector—an increase of almost 45 percent compared with business-as-usual. Financial Services can capitalize on AI technologies to relieve knowledge workers from mundane, repetitive tasks such as generic customer queries, mortgage reviews and market research. Overall, this sector will benefit from US$1.2 trillion in additional GVA in 2035. Even labor-intensive sectors— where productivity growth is notoriously slow— will experience significant increases in GVA growth rates. Education will

see a boost from 0.9 percent to 1.6 percent by 2035 and Social Services from 1.6 percent to 2.8 percent, yielding substantial increases in economic output (an additional US$109 billion and US$216 billion in GVA respectively).

Please see answers that follow for "disadvantages for consumers and for businesses of utilizing products or services facilitated by algorithms, artificial intelligence, or predictive analytics."

………

From a technical perspective, is it sometimes impossible to ascertain the basis for a result produced by these technologies? If so, what concerns does this raise?

………

One of the things that concerns stakeholders about Artificial Intelligence is the idea that our world will be filled with AI "black boxes" that will make important decisions without accountability or recourse. Industry is positioned and ready to be a leader and partner with government to ensure there is a pathway for Responsible AI. We define Responsible AI as the practice of using AI with good intention to empower employees and businesses, and fairly impact consumers and society – allowing companies to engender trust and scale AI with confidence. It takes into consideration key issues such as: privacy, liability, transparency, bias, human values, data protection, and cybersecurity. For the purpose of ascertaining results, industry and government should work together to build and support technologies rooted in the notion of TRUST:

T	**R**	**U**	**S**	**T**
Trustworthy Safe, honest, and diverse in perspectives	**Reliable** Enabling enhanced judgement and making better decisions	**Understandable** Interpretable and transparent decision-making	**Secure** Privacy of company and customer information and data	**Teachable** Human-centric design, aiming for humans + machines to co-create, inform, and educate each other

Accenture's core Responsible AI principles are a code of ethics and business values around building:

TRUSTWORTHY AI that is safe, honest, and diverse in perspectives, thereby earning the trust of consumers, and society. AI systems need to be trusted agents of their creators, vested with the same duties for responsible and effective behavior as any traditional software system or person would be.

RELIABLE AI that enables enhanced judgement and makes better decisions based on a diverse set of key values incorporated into algorithms. Reliable systems act consistently and are sufficiently robust to perturbations – so new observations are not going to make the system act erratically or against the organization's values.

UNDERSTANDABLE AI that allows for transparency and interpretability in decision-making. It enables both a non-technical person (e.g. organization executive or customer) and a technical expert (e.g. an AI/machine learning expert) to gain insight into how an algorithm works and why it made a given decision. There is an important difference between merely meeting any legal requirements to be transparent versus a desire to establish trust and prioritize understandability. Interpretability tends to be the focus of what organizations such as DARPA call "explainable AI". DARPA defines explainability as the ability of machines to: 1) explain their rationale 2) characterize the strengths and weaknesses of their decision-making process and 3) convey a sense of how they will behave in the future.
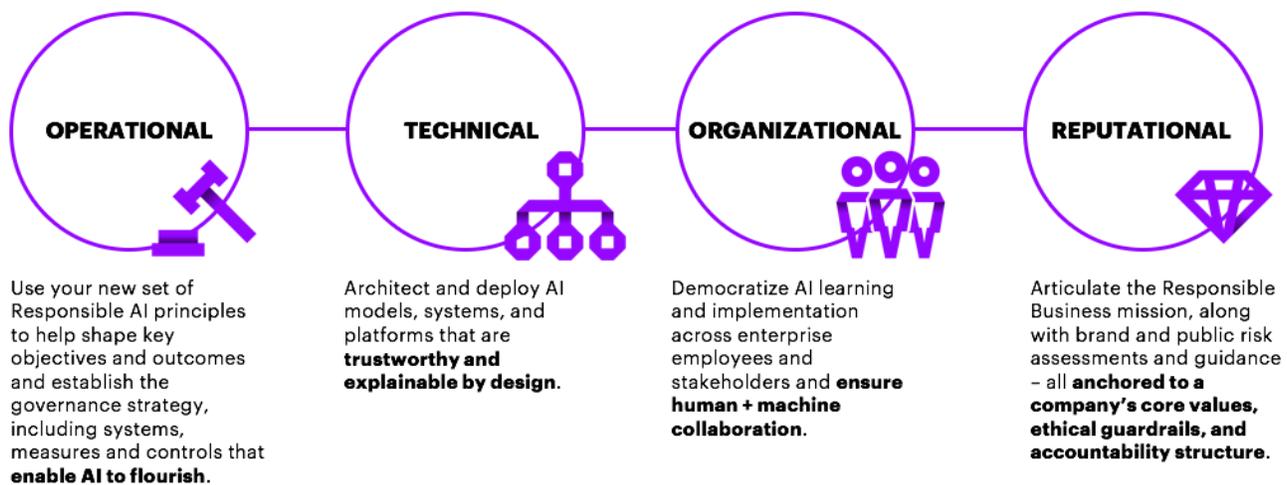
In addition, organizations may want to consider how to proactively *justify their design choices* by explaining:

- why they chose a particular data set to draw inferences,

- why these inferences are relevant (and ethical*) for the chosen decision they are trying to make, and

- whether the data and methods used to draw the inferences are accurate and statistically reliable for the population they are trying to serve. A data set full of Iowans would not serve the population of California whose population is full of different characteristics.

SECURE AI that maintains the privacy of company and customer information and data. The security of AI is intended to ensure that sensitive customer information is inaccessible to bad actors and that AI systems are not vulnerable to manipulation.

TEACHABLE AI with a human-centric design, aiming for humans + machines to co-create, inform, and educate one another. In considering the human impact of integrating AI into high risk areas, such as criminal justice and health care, organizations can set out to design, build, and deploy AI systems whereby human responsibility is enhanced. Whereas transparency provides insights into the systems driving decision-making, teachable AI enables the ability to change or alter how consumers interact with that system. An added complexity, that is aligned with the issues of transparency, is accountability of human and algorithmic systems. Presumably we build algorithms, at least in part, to standardize and remove human bias even while we simultaneously say AI can biased and call for human oversight. Because of this, we need to remain focused on tracing decision making, not just of algorithms but of people.

OPERATIONAL	TECHNICAL	ORGANIZATIONAL	REPUTATIONAL
Use your new set of Responsible AI principles to help shape key objectives and outcomes and establish the governance strategy, including systems, measures and controls that **enable AI to flourish**.	Architect and deploy AI models, systems, and platforms that are **trustworthy and explainable by design.**	Democratize AI learning and implementation across enterprise employees and stakeholders and **ensure human + machine collaboration.**	Articulate the Responsible Business mission, along with brand and public risk assessments and guidance – all **anchored to a company's core values, ethical guardrails, and accountability structure.**

In order to successfully develop and execute an internal governance model, Accenture recommends the following four pillars:

- **Operational**: Without effective oversight, ethical considerations can fall by the wayside, valid concerns can be disregarded, and opportunities for responsibility can be lost. Companies can define their own core principles and establish a governance strategy with systems, measures, and controls that enable AI to flourish – without sacrificing ethics. The governance pillar aims to support innovation through responsible action.
- **Technical:** The traditional area of focus, the technical pillar aims to define Data Science-based models, systems, platforms and tools that can help address concerns around Responsible AI through trustworthy and explainable design. One solution is the Accenture Fairness Tool, which is designed to understand and enable humans to better visualize and understand how bias may emanate in both the data and the algorithmic models that are at the core of AI systems. The tool allows for non-technical and technical people in an organization to have a conversation

around whether the models are aligned with their core values and, most important, enables better decision-making -- rather than making decisions autonomously.

- **Organizational:** As with any high-level initiative, success or failure of Responsible AI rests on the people who strive to achieve it. The organizational pillar focuses on democratizing AI learning and implementation to ensure human and machine collaboration. This includes the democratization around skills and training needed, how people should be arranged within an organization, and what operational processes best support success.
- **Reputational**: A key source of value for any organization, this pillar aims to define the impact of Responsible AI on an organization's brand. From public scrutiny, to customer trust and the existential impact of fairness for high-scale tools, it includes ways to understand the effect that an organization's core values, ethical guardrails, and accountability structure will have on business.

……

What are the advantages and disadvantages of developing technologies for which the basis for the results can or cannot be determined? What criteria should determine when a "black box" system is acceptable, or when a result should be explainable?

……

While "black box" AI is clearly limited by its inability to explain its reasoning to human users, it can nonetheless function well in three types of applications:

- *Pattern recognition tasks* where the cost of failure is low.
- *"Closed-loop" systems* where a real-time response is critical, and/or the pace of decision-making is too fast to allow for human intervention, such as real-time online retail pricing, movie or music recommendations, or driverless cars.
- *Interactive response systems* such as robots and chatbots.

However, in many other applications—especially business situations such as financial services and health care, where an explanation of the underlying reasoning is critical for decision-makers—machine learning's lack of explainability is, rightfully, a major barrier to the adoption of AI. The underlying dynamic is that as the human impact and implications of AI decisions increase, so does the need to explain especially in high risk areas such as health care, financial services, and security. If AI lacks the ability to explain itself in these areas, then the risk of it making a wrong decision may outweigh the benefits it could bring in terms of the speed, accuracy and efficiency of decision-making. The effect may impact the ability of the industrial uptake of the technology.

In the Accenture Technology Vision 2018, we present our view of how technology will augment and enhance human skills to enable us all to listen more closely to customers and employees, connect to them on their own terms, and partner with them to achieve personal goals. By advancing AI for responsibility, fairness, and transparency, organizations can create a collaborative, powerful new member of the workforce. Explainable and more responsible AI will be the backbone of the intelligent systems of the future that enable the intelligent enterprise. In playing this role, Explainable AI won't replace people, but will complement and support them so they can make better, faster, more accurate and more consistent decisions.[xii]

Common Principles and Ethics in the Development and Use of Algorithms, Artificial Intelligence, and Predictive Analytics

---

……

What are the main ethical issues (e.g., susceptibility to bias) associated with these technologies? How are the relevant affected parties (e.g., technologists, the business community, government, consumer groups, etc.) proposing to address these ethical issues? What challenges might arise in addressing them?

……

AI has exciting implications in the consumer, healthcare, transport, and professional services industries. It promises to deliver learning and evolving systems that are more efficient, safer, smarter and better able to operate at scale. But these opportunities come with associated ethical challenges:

- *Decision making and liability* - As AI use increases, it will become more difficult to apportion responsibility for decisions. What steps need to be taken to prevent mistakes and manage risks? If mistakes are made which cause harm, who should bear the risk?
- *Transparency* - When complex machine learning systems have been used to make significant decisions, it may be difficult to unpick the causes behind a specific course of action. The clear explanation of machine reasoning is necessary to determine accountability.
- *Bias* - Machine learning systems can entrench existing bias in decision making systems. Care must be taken to ensure that AI evolves to be non-discriminatory.
- *Human values* – Humans must work alongside AI especially because AI systems have no default values or "common sense".
- *Data protection and IP* - The potential of AI is rooted in access to large data sets. What happens when an AI system is trained on one data set, then applies learnings to a new data set? Could this infringe data protection and privacy rights? Who, if anyone, should own the output of AI thought processes?
- *Cybersecurity* - As more reliance is placed on AI, the importance of cybersecurity will increase. Security must be a top priority for all actors if trust is to be maintained.[xiii]

What is Bias?

It is important to know that bias has two main types: experimental or data bias, and societal bias. Bias in data or the design of a model/experiment is often measurable. This exists because a dataset may be lacking enough women, or people of color, or Minnesotans over the age of 65. As this bias is measurable, it is possible to correct by adding additional data, or changing parameters. Societal bias, however, is not corrected by data. In fact, we can (and often do) have 'perfect' data that reflects an imperfect world. It is telling, for example, that one employer's AI tool for hiring learned to eliminate resumes if they were determined to be from women. This AI was built on 'perfect' data.

The Impact of Bias[xiv]

However objective we may intend our technology to be, it is ultimately influenced by the people who build it and the data that feeds it. Technologists do not define the objective functions behind AI independent of social context. Data is not objective, is it reflective of pre-existing social and cultural biases. In practice, AI can be a method of perpetuating bias, leading to unintended negative consequences and inequitable outcomes.

Today's conversation about unintended consequences and fair outcomes is not new. In 1971, the U.S. Supreme Court established the notion of "disparate impact" — the predominant legal theory used to review unintended discrimination. Specifically, the Griggs vs. Duke Power Company ruling stated that independent of intent, disparate and discriminatory outcomes for protected classes (in this case, with regard to hiring), were in violation of Title VII of the Civil Rights Act of 1964. Today, this ruling is widely used to evaluate hiring and housing decisions, and it is the legal basis for inquiry into the potential for AI discrimination. Specifically, it defines how to understand "unintended consequences" and whether a decision process's outcomes are fair. While regulation of AI is in early stages, fairness will be a key pillar of discerning and avoiding adverse impact.

The field of AI ethics draws an interdisciplinary group of lawyers, philosophers, social scientists, programmers, and others. Influenced by this community, Accenture Applied Intelligence has developed a fairness tool to understand and enable humans to address bias in both the data and the algorithmic models that are at the core of AI systems.

Our tool measures disparate impact and provides users options to correct these impacts for predictive parity to advance equal opportunity. The tool exposes potential disparate impact by investigating the data and model. The process integrates with the existing data science processes. Step 1 in the tool is used in the data investigation process. Steps 2 and 3 occur after a model has been developed. In its current form, the fairness evaluation tool works for classification models, which are used, for example, to determine whether to grant a loan to an applicant. Classification models group people or items by similar characteristics. The tool helps a user determine whether this grouping occurs in an unfair manner and provides methods of correction.

One priority in developing this tool was to align with the agile innovation process competitive organizations use today. Therefore, our tool needs to be able to handle large amounts of data, so it wouldn't keep organizations from scaling proof-of-concept AI projects. It also needs to be easily understandable by business leaders (not just AI or data experts). It also needs to operate alongside existing data science workflows, so the innovation process is not hindered.

Issues with Correcting Bias[xv]

Our tool does not simply dictate what is fair. Rather, it allows an organization to assess potential bias within the parameters set by its users who ultimately need to define sensitive variables, error terms and false positive rates. Their decisions should be governed by an organization's values and basic principles that an organization will follow when implementing AI to build trust with its stakeholders, avert risks to their business, and contribute value to society.

Success depends not just on offering solutions to improve algorithms, but also on its ability to explain and understand the outcomes. It is meant to facilitate a larger conversation among data scientists and non-data scientists. By creating a tool that prioritizes human engagement over automation in human-machine collaboration, we aim to inspire the continuation of the fairness debate into actionable ethical practices in AI development.

……

Are there ethical concerns raised by these technologies that are not also raised by traditional computer programming techniques or by human decision-making? Are the concerns raised by these technologies greater or less than those of traditional computer programming or human decision-making? Why or why not?

……

There have been many examples in the media about incidents where one must consider the impact that AI can have on trust:

- *Immediate:* Deployment enabled by cloud computing is integrated into existing systems and customer services without delay. Agile prototyping means that minimum viable products and improvements are constantly pushed out.
- *Impactful*: AI launches to all customers globally. Level of intimacy with customers has an expected reciprocity. Imagine when AI intersects with very personal decisions around employment, housing, health, education, lending, voting.
- *Invisible*: Due to hyper-personalization and customization of AI, it is impossible to tell if one user experience involving AI is a result of model bias. Would need to perform meta-analyses. Customers will understand at some point but will not always register until the very end (sometimes with the aid of a physical manifestation).

For this reason, internal AI governance and transparency should continue to remain front and center for organizations deploying the technology.

……

Is industry self-regulation and government enforcement of existing laws sufficient to address concerns, or are new laws or regulations necessary?

……

AI is a rapidly maturing and fast changing field in which new, transformative technologies are emerging almost every week. Government should avoid regulating the technology per se, on the basis that any regulation of that kind will become obsolete

as the technology evolves rapidly. Instead, our view is that government has a role to work with industry and other AI stakeholders in order to address the issues that arise in specific uses and applications of AI that could have certain consequences that impact human life and public safety.

In addition, there are roles for government to play, working with industry, that include data protection, ethics and transparency, IP ownership, risk allocation and cybersecurity. Consideration should also be given to global, industry driven guidelines, standards and best practices that can help create and safeguard trust at the heart of AI-driven systems and business models and permit the flexibility for innovation, allowing codes to develop with the technology. Examples include, the IEEE Global Initiative for Ethical Considerations in AI and AS and its nine pipeline standards on Ethically Aligned Design.

These multi-stakeholder initiatives have the strongest influence on creating industry cross-fertilization and equal access on AI resources for entrepreneurs and big players alike. They also help to identify gaps in existing standards and certifications, which ecosystem players can then act upon. A one-size fits-all solution should be avoided as applications vary across sectors and industries.[xvi]

……

Should ethical guidelines and common principles be tailored to the type of technology involved, or should the goal be to develop one overarching set of best practices?

……

Any ethical guidelines, common principles, or best practices should be derived from global, industry driven guidelines, standards and best practices that can help create and safeguard trust at the heart of AI-driven systems and business models and permit the flexibility for innovation, allowing codes to develop with the technology.

Consumer Protection Issues Related to Algorithms, Artificial Intelligence, and Predictive Analytics

---

……

What are the main consumer protection issues raised by algorithms, artificial intelligence, and predictive analytics?

……

Predictive and prescriptive analytics are not problematic and have a wide range of use cases that are very beneficial to society. The issues become whether these technologies create the idea that certain characteristics are immutable (it will happen, it must happen vs. it may happen, it is likely to happen). Responsible AI systems should incorporate constructive dissent – the ability for consumers to correct, delete, remove and ask for an explanation in relation to personal decisions relating to housing, education, loans, employment, and other high impact areas so as to avoid algorithmic determinism, or the concept that the combination of measurement bias and feedback loops can create a self-reinforcing bias with a wide range of detrimental effects.

Another concern is algorithmic determinism, which was presented at the FTC hearings by Accenture's Global Lead for Responsible AI, Dr. Rumman Chowdhury. Algorithms are designed to distinguish patterns and prescribe actions to fulfill that predicted pattern. In doing so, they may create self-fulfilling prophecies where individuals act as per prediction because their experience or treatment is shaped by that assumption.

There are two potential harms: manipulation, and the limitation of opportunity. We observe manipulation in the filter bubble, where recommendation algorithms designed by media outlets are optimized to maximize the number of clicks. In doing so, they promote questionable and/or bias-reinforcing media, which creates in the viewer a false understanding of issues and a lack of knowledge about alternative viewpoints. The limitation of opportunity can result if we categorize an individual as being of a type (e.g. 'low performing customer' or 'high performing customer') and then target or treat them accordingly. If we provide sub-par service or offers to a predicted 'low performer,' they are more likely to respond according to the environment that is shaped for them. Extend this to risk assessments used in policing, or social credit scores, or hiring, and

there is a strong possibility that existing biases may reinforce and cement existing socioeconomic, racial and gender biases with digital disparate treatment.

Please see also previous answer above for the following question:
> "*From a technical perspective, is it sometimes impossible to ascertain the basis for a result produced by these technologies? If so, what concerns does this raise?*

………

In recent years, the FTC has held public forums to examine the consumer protection questions raised by artificial intelligence as used in certain contexts (e.g., the <u>2017 FinTech Forum</u> on artificial intelligence and blockchain and the <u>2011 Face Facts Forum</u> on facial recognition technology). Since those events, have technological advancements, or the increased prevalence of certain technologies, raised new or increased consumer protection concerns?

………

These digital technologies have been transforming the economy in such a way that companies have found that they are no longer just a retail company or an agricultural company – but also a technology or a technology-empowered company. Digital technologies, now powered by AI, enable companies to see and understand their customers in new and exciting ways, thereby allowing companies to move into new business models and arenas with ease. These "new" entrants must be introduced to concepts that are fundamental to consumer protection but a matter of first impression to these companies. Imagine for example a company that awards its clients digital rewards on the basis of certain behaviors and characteristics but does not realize the possibility of triggering the Fair Credit Reporting Act. In some respects, these technologies have not raised new or increased consumer protection concerns per se, but rather amplify the audience to those who need to understand it for the first time.

The democratization of AI has been, in part, one of the reasons for its quick adoption, but also the reason for misuse. As AI becomes a standard tool for non-technology uses, new situations arise that require thoughtful consideration with regard to citizen protection. A few examples include:

- Denial of service based on tracking of consumer's daily lives through IoT.
- Deepfake technology which may impersonate an unwilling and unwitting individual with no way to verify its veracity.
- Abuse of domestic IoT devices for purpose of stalking or harassing.

………

What roles should explainability, risk management, and human control play in the implementation of these technologies?

………

Please also see previous answers for the following questions:
- "*From a technical perspective, is it sometimes impossible to ascertain the basis for a result produced by these technologies? If so, what concerns does this raise?*
- *What are the advantages and disadvantages of developing technologies for which the basis for the results can or cannot be determined? What criteria should determine when a "black box" system is acceptable, or when a result should be explainable?*

………

What educational role should the FTC play with respect to these technologies? What would be most useful to consumers?

………

We believe that efforts should be made to improve understanding of, and engagement with, AI. In the last year, there has been an ever-increasing velocity of articles, blogs, speeches and thinking raising concerns about AI. The reality is that, in some instances, AI itself does not introduce novel questions or concerns (for example, the classic *trolley* question that is being raised in the context of autonomous vehicles). Yet, the hype around the technology has the potential to create misguided understandings of both the promise and peril of AI. The focus on the vivid risks rather than the potential benefit is negative for the public overall as it potentially impedes adoption of the technology and thereby delays the uptake of the benefits brought by the technology.

Given the public's increased interest in this area, this is the right time for educational campaigns around: (1) topics like data privacy and cybersecurity and (2) helping consumers and citizens understand their rights and increase their digital fluency in a world of AI. This level of tangible engagement will enable an informed dialogue within our communities, accountability between public and private sectors as well as a drive to innovate.

Competition Issues Related to Algorithms, Artificial Intelligence, and Predictive Analytics

---

....

Is the current antitrust framework for analyzing mergers and conduct sufficient to address any competition issues that are associated with the use of these technologies? If not, why not, and how should the current legal framework be modified?

....

It is important that any new regulations or laws pertaining to AI do not unintentionally stifle development of this ground-breaking technology. According to our research, AI will significantly boost labor productivity in developed economies and add trillions in economic output in the coming decades. Looking at 12 of the world's largest economies, Accenture believes AI has the potential to double annual economic growth rates by 2035 by changing the nature of work and creating a new relationship between man and machine.[xvii]

Other Policy Questions

....

How quickly are these technologies advancing? What are the implications of that pace of technological development from a policy perspective?

....

The FTC rightly asks what is the current state of AI, because there is a high degree of hype in the media today. We see discussions ranging from AI becoming spontaneously evil weaponry to AI solving all the world's problems. In many respects, the current state of AI is that we are not "there" yet; AI needs human helpers and developers to execute its tasks with excellence. AI can undertake tasks but not necessarily end to end processes of any complexity. However, even in relatively "weak AI" there are many things that AI can do very well, especially where tasks are (1) repeatable and/or (2) require time-consuming analysis of large datasets.

We are seeing that the technology is progressing rapidly, and we believe that the entirety of the technology market will be impacted by the emergence of AI. Accenture research on the impact of AI in 12 developed economies reveals that AI could double annual economic growth rates by2035 by changing the nature of work and creating a new relationship between man and machine. The impact of AI technologies on business is projected to increase labor productivity by up to 40 percent and enable people to make more efficient use of their time.[xviii]

At the same time, companies should take into consideration the global dialogue occurring in this arena, pursue responsible data practices and seek to deepen partnering with regulatory authorities.

....

How can regulators meet legitimate regulatory goals that may be raised in connection with these technologies without unduly hindering competition or innovation?

....

Please see previous answer: *How quickly are these technologies advancing? What are the implications of that pace of technological development from a policy perspective?*

[i] **AI Explained: A Guide for Executives, Accenture:** https://view.pagetiger.com/AI-Explained-A-Guide-for-Executives/2018/160x160AIbrochurever10.pdf

[ii] Ibid.

[iii] Ibid.

[iv] **Realizing the Promise of Intelligent Automation, Accenture:** https://www.accenture.com/t20161124T183925Z__w__/us-en/_acnmedia/PDF-36/Accenture-Realizing-Promise-Intelligent-Automation.pdf#zoom=50

[v] **How AI Boosts Industry Profits and Innovation, Accenture**: https://www.accenture.com/t20171005T065828Z__w__/us-en/_acnmedia/Accenture/next-gen-5/insight-ai-industry-growth/pdf/Accenture-AI-Industry-Growth-Full-Report.pdfla=en?la=en

[vi] **Accenture response to "Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan"** : https://www.nitrd.gov/rfi/ai/2018/AI-RFI-Response-2018-Adelina-Cooke-Accenture.pdf

[vii] **AI Momentum, Maturity, and Models for Success; SAS, Accenture Applied Intelligence, and Intel with Forbes Insights**: https://www.accenture.com/t20180919T202227Z__w__/us-en/_acnmedia/PDF-86/Accenture-AI-Momentum-Final.pdf#zoom=50

[viii] Ibid.

[ix] Ibid.

[x] Ibid.

[xi] **As companies embrace AI, it's a job-seeker's market, Reuters:** https://www.reuters.com/article/us-usa-economy-artificialintelligence/as-companies-embrace-ai-its-a-job-seekers-market-idUSKCN1MP10D

[xii] **Explainable AI, Human-Machine Collaboration, Accenture:** https://www.accenture.com/us-en/insights/technology/explainable-ai-human-machine

[xiii] **Responsible AI and Robotics: An Ethical Framework, Accenture:** https://www.accenture.com/gb-en/company-responsible-ai-robotics

[xiv] **Auditing Algorithms for Bias (Harvard Business Review), Accenture:** https://hbr.org/2018/10/auditing-algorithms-for-bias

[xv] Ibid.

[xvi] **Engage Digital, Accenture:** https://www.accenture.com/gb-en/insight-engage-digital

[xvii] **Artificial Intelligence is the Future of Growth, Accenture:** https://www.accenture.com/us-en/insight-artificial-intelligence-future-growth

[xviii] **Artificial Intelligence is the Future of Growth, Accenture:** https://www.accenture.com/us-en/insight-artificial-intelligence-future-growth