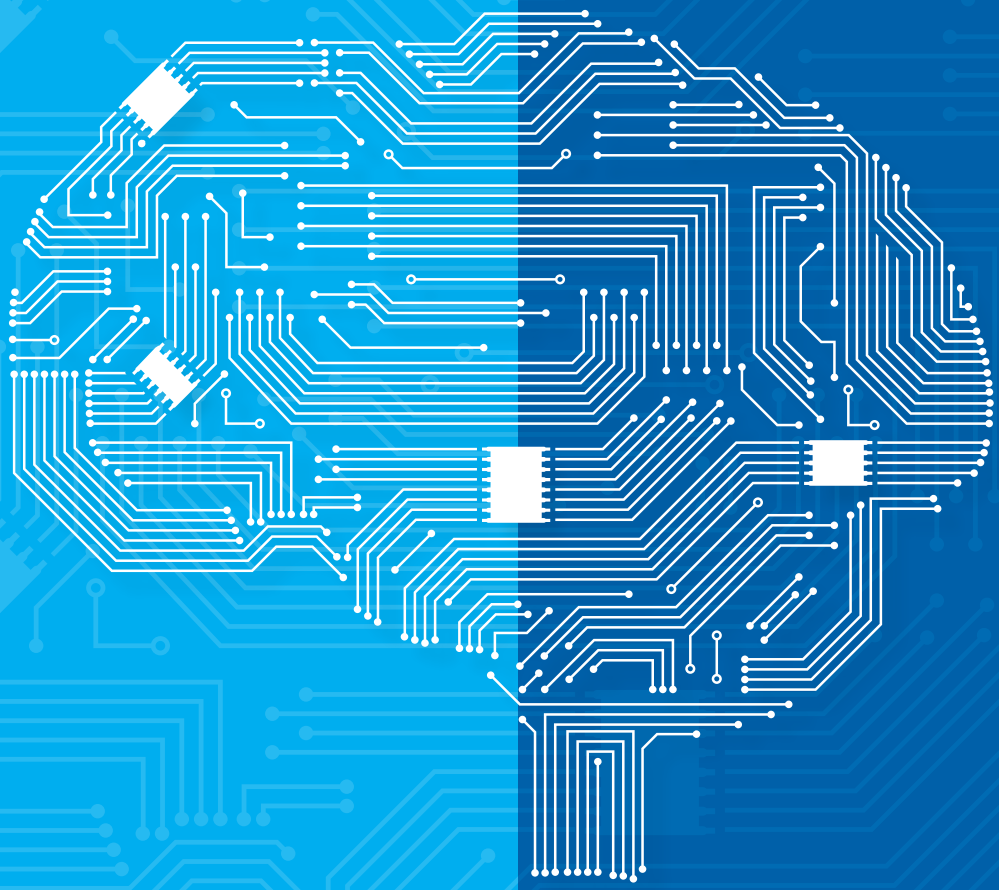

2019 EDELMAN AI SURVEY

SURVEY OF TECHNOLOGY EXECUTIVES AND THE GENERAL POPULATION SHOWS EXCITEMENT AND CURIOSITY YET UNCERTAINTY AND WORRIES THAT ARTIFICIAL INTELLIGENCE COULD BE A TOOL OF DIVISION

March 2019

Edelman

Contents

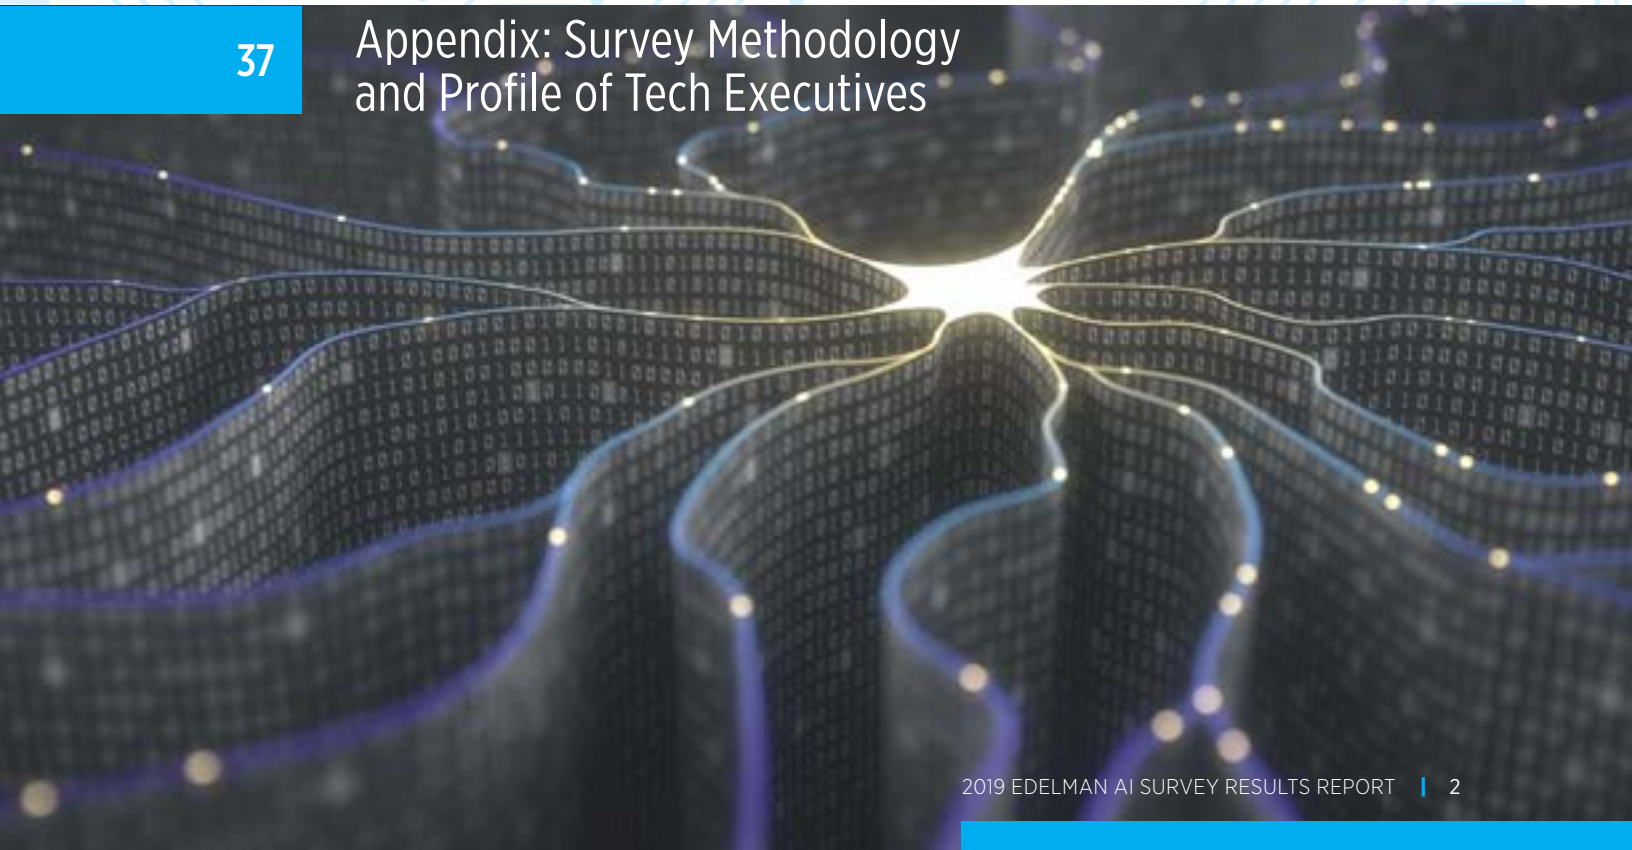

Executive Summary

Born in the 1950s, artificial intelligence (AI) is hardly new. After suffering an "AI Winter"* in the late 1980s, recent advances with more powerful computers, more intelligent software and vast amounts of "big data" have led to breakneck advances over the last several years mostly based on the "deep learning" breakthrough in 2012. Every day, there are headlines extolling the latest AI-powered capability ranging from dramatic improvements in medical diagnostics to agriculture, earthquake prediction, endangered wildlife protection and many more applications. Nevertheless, there are many voices warning about a runaway technology that could eliminate jobs and pose an existential threat to humanity.

Given the dramatic advances and dichotomy of views, AI is a widely discussed topic in opinion pieces and news stories, conference keynotes, various surveys, and is a recurrent theme in television and movies. The Edelman AI Center of Expertise and Edelman Intelligence business units wanted to probe deeper. Specifically, we wanted to determine current viewpoints but also to compare perceptions of the public with those of people who work most closely with technology. We called this latter group the "tech execs" as they work daily in development or deployment with leading technologies. More than most, these tech execs have a front row seat on what is happening with AI. We found that nearly 9 in 10 survey respondents from both groups believe AI is transformational, and that it is leading to a technological revolution. Both groups are generally optimistic about the long-term impact of AI though the results clearly show the tech execs are more so. This optimism delta between the groups was consistent across almost every question. Responses clearly show both groups expect AI to provide business benefits across every industry.

However, the survey also reveals some very real and troubling concerns. Majorities of both groups believe that initially, AI will eliminate more jobs than it creates and lead to greater income inequality. Strong majorities believe that the wealthy will benefit from AI while nearly half expect the poor will be harmed. Approximately 80 percent of respondents expect AI to invoke a reactionary response from those who feel threatened by the technology. Additionally, there are also worries about the dark side of AI, including concerns by nearly 70 percent about the potential loss of human intellectual capabilities as AI-powered applications increasingly make decisions for us. Furthermore, 7 in 10 are concerned about growing social isolation from an increased reliance on smart devices.

The survey reveals the many positive benefits but also the potential that AI can be a powerful tool of division. Nearly 70 percent believe AI interjects greater possibilities for digitally enhanced "group think," lessening creativity and freedom of thought. AI thrives on a lot of centralized information. As autocratic powers are able to access and leverage that content, they could increasingly control the levers of disinformation thereby further manipulating their populations. Adding to that, nearly half expect AI-generated 'deepfake' audio and visuals will further erode public trust in what is real, and a third believe these could lead to a war.

Despite these reservations, what comes through from the responses to questions and from verbatim quotes is an underlying curiosity – a desire to know where AI will lead us. Given a choice of words to describe how they feel about AI, "curiosity" wins out over "fear," and this is what drives us forward despite the apparent risks. It could be that Google CEO Sundar Pichai is right when he said at the 2018 World Economic Forum in Davos: "AI is probably the most important thing humanity has ever worked on. I think of it as something more profound than electricity or fire." Added to this, Kai-Fu Lee – CEO of Chinese venture capital firm Sinovation Ventures, author of the "AI Superpowers: China, Silicon Valley and the New World Order" and dubbed the "Oracle of AI" by 60 Minutes – claimed in a January 2019 interview: "I believe [AI] is going to change the world more than anything in the history of mankind. More than electricity."

* The term "AI Winter" first appeared in 1984 as the topic of a public debate at the annual meeting of AAAI. Two leading AI researchers warned the business community that enthusiasm for AI had spiraled out of control in the 1980s and that disappointment would certainly follow. Three years later, the billion-dollar AI industry began to collapse.

Technology

"Artificial intelligence is helping to pave the way for future generations.

– 32-year-old male, general population

Revolutionary Impact of AI

Though some may think of robots taking over the world when they envision artificial intelligence, AI is a broad term for many underlying fields such as natural language processing and image recognition. There is no consensus definition for AI, but we can say it is a collection of digital tools that enable machines to perceive, learn and make decisions like humans. These tools are advancing at a rapid, possibly exponential rate. Most of the advances so far use a subdiscipline of AI known as machine learning that is based on mathematical algorithms, whose performance improves as a function of processing more data. There are thousands of such algorithms already in use today. Deep learning is a type of machine learning designed to mimic neurons in the human brain. The most significant advances so far have been through the several versions of deep learning neural networks. There are still other forms of AI that remain in early stages of development but could offer their own substantial benefits.

Nearly every day brings news of additional technical breakthroughs and useful AI applications. These range from more accurate medical diagnostic tools and personalized drug treatment to boosting the amount of information carried over fiber-optic telecommunication networks, broad uses in business including spotting fraud in expense reports, improving agriculture yields, offering empirical proof that fibromyalgia is a real disease, improving education, minimizing the impact of natural disasters and even helping with methods to identify potential school shooters.

Still, machine learning and deep learning are part of what is known as "narrow AI." They essentially perform one task, albeit very well and with a certain degree of intelligence, within one field such as voice or image recognition. General AI—also known as "strong AI"—is far more sophisticated and does not yet exist. This type of AI would be able to perform any general task asked of it, much like a human. Ultimately, a general AI — with advanced cognitive abilities, understanding of its environment and ability to process vast amounts of information at computer speed — could lead to the "singularity," a super-intelligent machine that surpasses even the smartest human. Though it is possible this super intelligence might never be realized, futurists envision the singularity could be achieved in the next 20-30 years.

AI is fundamentally different from previous technical advances. The development of AI is creating a Cognitive Era, with machines able to perform tasks beyond the capabilities of people. Yet, for all the apparent gains so far, AI is likely still in its infancy. The full impact of AI technologies on business, the global economy, individual people and all of society are unknown. Almost assuredly these impacts will be profound, including a possible human-species-defining moment.

AI Leads in Current Technology Significance

Much has been said in various forums about AI being the most significant technology of this time. Our survey confirmed this perception as nearly 7 in 10 tech execs ranked AI in the top 3 of most significant technologies over the next 5-10 years. This was followed by cloud computing with 58 percent and big data/internet of things at 50 percent. On the other end of the significance spectrum are augmented and virtual reality with only 18 and 16 percent, respectively, placing these in their top 3.

RANKED #1-3 FOR SIGNIFICANCE	TECH EXECS
Artificial Intelligence	69%
Cloud Computing	58%
Big Data / Internet of Things	50%
Robotics	26%
Blockchain	24%
Augmented Reality	18%
Nanotechnology	17%
Virtual Reality	16%
Human / Machine Integration	11%
Ambient Computing	10%

At the Center of a Technology Revolution

More than 9 in 10 of tech execs somewhat or strongly agree that AI is the next technological revolution and both groups believe AI-powered machines will take on mundane tasks, allowing people to have greater freedom to pursue more creative work and play.

"I will have it do my work and I can get the stuff done on a more creative side.

– 39-year-old female tech exec

The U.S. Leads in AI, for Now

Tech execs generally believe that the U.S. is at the forefront of AI innovation though the general population sees a close international race for leadership. Regardless of which perception is closer to reality, there's no doubt many countries are in hot pursuit of AI. A recent study titled Strategic Competition in an Era of Artificial Intelligence, suggests that the U.S. has been flat-footed in the global AI race.

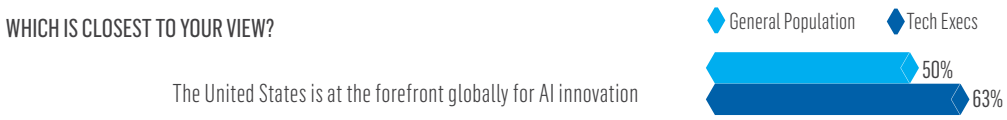

AI Leadership and Economic Future

Type "AI is reshaping the future" into a search engine and page one returns include: "How AI is reshaping the future of cybersecurity," "AI is reshaping the future of financial services," "How AI, machine learning and analytics are reshaping the future of IT Operations," "AI is reshaping the future of business," "5 ways AI is reshaping the future of e-commerce," "AI is reshaping the future of medicine," and "How AI is reshaping the consumer goods industry." Page two is full of the same, including how "How AI, drones and big data are reshaping the future of warfare."

Just how important is AI? Microsoft's Chief Envisioning Officer, Dave Coplin said recently that AI is "the most important technology that anybody on the planet is working on today." Management consulting group McKinsey forecasts that AI will add $13 trillion to the global economy by 2030. The Atlantic Council claims in a report that AI, and more specifically who will lead in AI, is a "Sputnik" moment. The reference is to the Russian satellite launched in 1957 that marked the beginnings of a space race which ultimately led to the U.S. being first to send men to the moon. The national commitment to this endeavor led to technologies that provided a significant advantage for the U.S., not only in space but in computing and other fields.

The Atlantic Council report claims the U.S. is missing a similar moment, though this time it is China that may take the lead. News reports voice similar concerns, citing the lack of a coherent U.S. policy to foster AI innovation and the exact opposite from China. Eric Schmidt, CEO of Google parent company Alphabet said recently: "It's pretty simple. By 2020, [China] will have caught up. By 2025, they will be better than us. By 2030, they will dominate the industries of AI."

Our survey groups reported considerably different views on AI leadership. By a wide margin – 63 percent to 37 percent – tech execs believe the U.S. is at the forefront of global AI innovation. However, the general population sees this as a tight race with only 50 percent of respondents saying the U.S. has a clear lead.

General Population Understands AI Better Than Assumed

Looking at both of our sample groups, when asked which of the following technologies fall under the umbrella of artificial intelligence, both the general population and tech execs agreed on the top three of machine learning, robotics and machine automation. This could indicate the general population is better educated on AI than is generally assumed. Neural networks are used in an advanced form of machine learning known as deep learning. It's odd that 78 percent of tech execs would correctly select machine learning but only 41 percent selected neural networks This could mean that even for tech execs, their knowledge of AI may be limited.

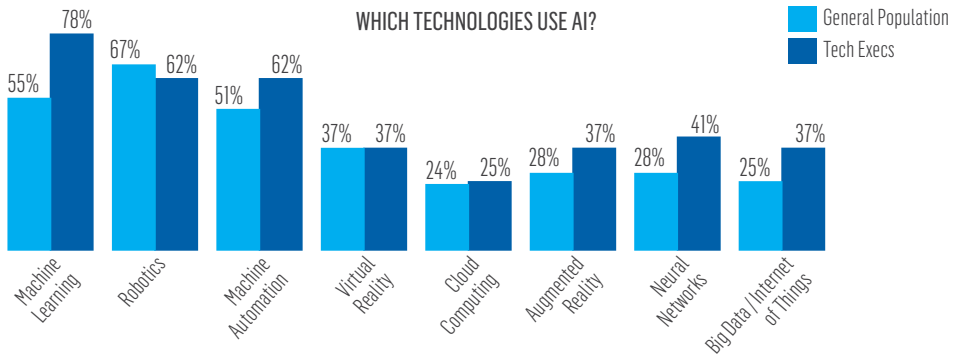

"AI is the creation of a computer program that can learn to think and function on its own, kind of like robots that don't need to be told what to do all the time.

– 29-year-old female, general population

But There Is Less Awareness for How AI Is Used

Both groups correctly recognized that voice assistants such as Apple Siri utilize artificial intelligence technologies. However, results were considerably lower in both groups for natural language processing, the key AI technology for those voice assistants. In the same way, both groups recognized that robots use AI, though responses were substantially lower for ingredient AI technologies such as computer vision. This suggests that many people in both groups understand that AI is currently used in some form within many consumer products, but there's less knowledge of how it works.

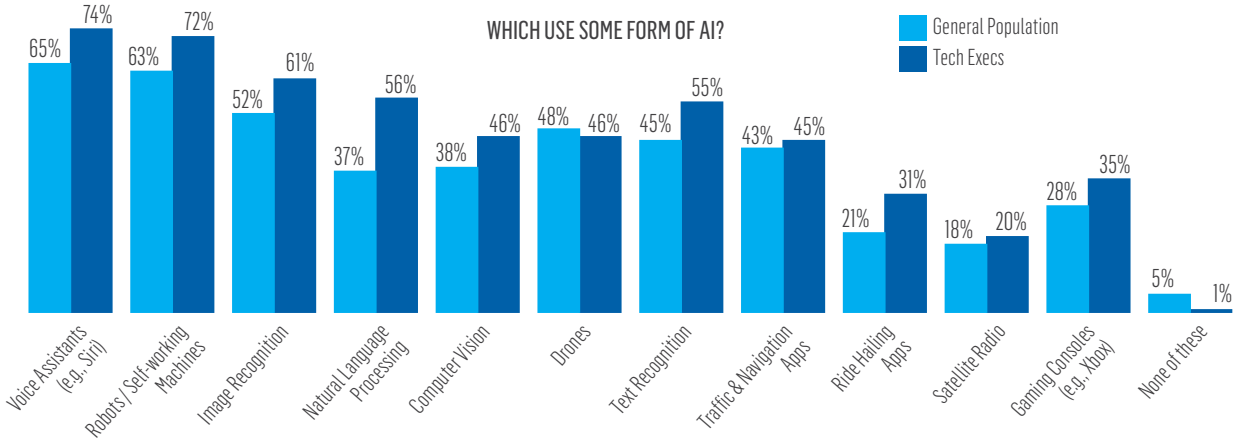

How Respondents Define AI

GENERAL POPULATION

"Artificial intelligence is machine behavior that emulates human thought.

– 52-year-old male

"Intelligence developed and generated by machines or computer.

– 66-year-old male

"Using a robot to do work that humans usually do.

– 54-year-old female

"A machine to interpret data and formulate conclusions or solutions.

– 25-year-old female

"A computer running itself, you hit go and it acts 'human.'

– 37-year-old female

TECH EXECS

> *AI is using computer algorithms to mimic human decision-making.*
>
> – 48-year-old female

> *Simulating consciousness electronically.*
>
> – 39-year-old male

> *Using a computer to do what normally the human brain does.*
>
> – 45-year-old female

> *Computers doing the work of the human brain.*
>
> – 63-year-old male

> *Computers performing tasks using programming that mimics human intelligence.*
>
> – 34-year-old female

AI Will Lead to Far-Ranging Improvements

Both groups strongly believe that AI is helpful across a variety of applications, especially with smart homes, personalized medicine, manufacturing improvements, creating smart toys and developing home assistive technologies for the elderly.

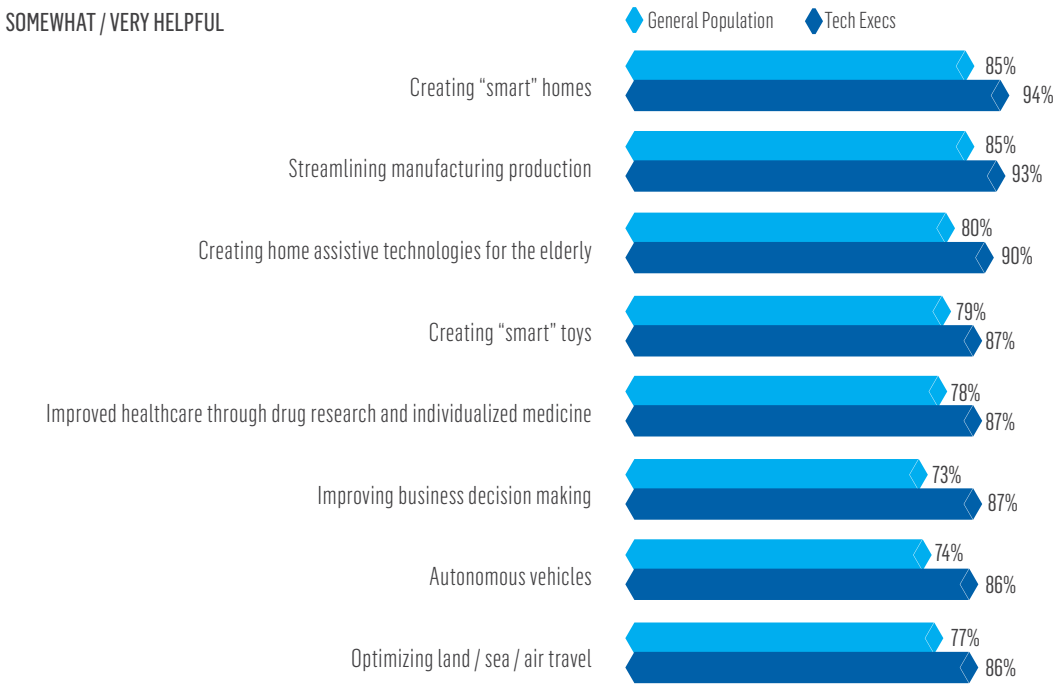

The Singularity Is Closer Than Thought

The singularity – the point at which AI-assisted machines surpass human intelligence – is predicted by visionaries such as Ray Kurzweil to arrive by 2045. MIT's Patrick Winston puts the date at 2040. After explaining the singularity, survey respondents are quite positive the singularity will be achieved even sooner, with 73 percent of tech execs saying this moment will arrive within 10 years, and nearly half of tech execs believe this will occur within 5 years. Perhaps the earlier prediction is a result of the impressive pace of technology development that could cause people to overestimate technological capabilities or achievements. Or it could be, as Professor Murray Shanahan believes, that "the human brain seems to be hard-wired to see intelligence where there is none: to see faces in clouds, to imagine a teddy bear as alive. Perhaps this is why some people so readily assume that human-level artificial intelligence will soon be with us."

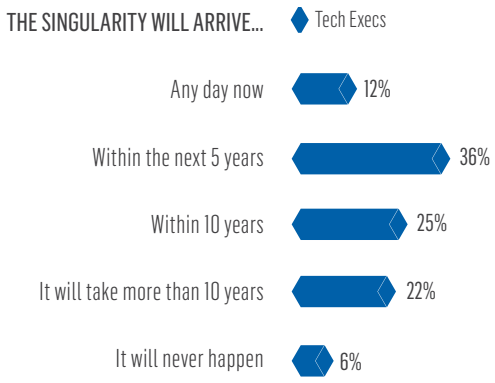

> *It's inevitable. I'm not particularly scared or worried by it.*
>
> – 32-year-old female tech exec

Singularity Concerns

While most people believe the singularity could arrive within 5-10 years, there's quite a bit of associated unease. When asked how the thought of machine superintelligence made them feel, 62 percent of the general population expressed negative emotions. Tech execs are evenly split on their feelings about superintelligence.

> *It makes me feel scared and excited at the same time for things to come.*
>
> – 49-year-old female, general population

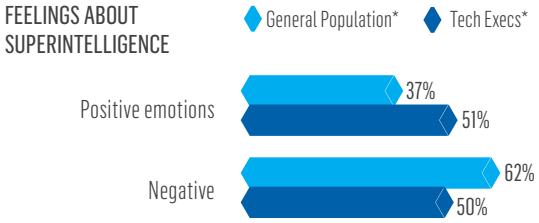

* Due to the open-ended nature of the question, results do not add up to 100%.

Split Views on the Singularity

The singularity is described as the point at which advances in artificial intelligence lead to the creation of a machine smarter than humans. Several notable futurists predict this will occur between 2030 and 2045. Not everyone agrees with this timeline. Estimates are based upon many factors, not the least of which is the continued increase in computing performance, eventually reaching as many connections as the human brain but operating much faster. Nick Bostrom and Vincent Muller conducted a survey of AI experts and found the median estimate is a one in two chance that high-level machine intelligence will be developed between 2040-2050, with a nine in ten chance by 2075. These experts believe that superintelligence will then be developed within the following 30 years, pushing the date to somewhere between 2070 and 2110. With the singularity realized, an AI-enabled machine will have achieved cognitive abilities beyond that of the most gifted human minds.

Not surprisingly, there are many views of whether this development will be positive or not. Computer scientist Stuart Russell has said that success creating human-level AI would be the biggest event in human history. Elon Musk has famously discussed how this superintelligence poses an existential threat to humanity. Musk is not alone in this concern. Bostrom considers the singularity potentially catastrophic for humanity. Nevertheless, it is quite possible that such superintelligence could be the key to addressing the complex assortment of human problems such as climate change. On the more optimistic side, futurist Ray Kurzweil acknowledges there will be a few bumps along the way, but that ultimately, superintelligence would allow for a symbiosis of man and machine.

Those bumps along the way could feel like mountains at the time they occur, such as AI-induced job displacement. Kurzweil and others believe that jobs will be lost, but newer ones will be created. Ultimately, his perspective is that AI and superintelligence will enhance humanity.

No one knows with any reasonable certainty when superintelligence will arrive. Professor Murray Shanahan at London's Imperial College believes the media often gives the impression that human-level AI of the sort we see in sci-fi movies is just around the corner. But it's almost certainly decades away. However, survey respondents believe the singularity will be achieved much sooner than do the experts, with 61 percent of the general population and 73 percent of tech execs saying this moment will arrive within 10 years. Views on whether this will be a positive development are quite split. About half of tech execs hold a positive view while just over a third of the general population expressed positive views.

Feelings About Superintelligence Possibilities

Survey respondents reflect a broad range of views about superintelligence with a seeming balance between those who are excited and optimistic and those scared or worried. In general, tech execs are more optimistic than is the general population.

GENERAL POPULATION

It scares me. Why would we create something that is smarter than us? It is not smart at all.

– 50-year-old female

I'm curious because I believe it part of natural evolution.

– 40-year-old male

Considering the state of the world I would welcome the AI overlords at this point.

– 34-year-old female

Apprehensive, but curious and hopeful that it benefits all life.

– 27-year-old male

Makes me feel excited.

– 21-year-old female

TECH EXECS

Great opportunity and also great potential for harm.

– 48-year-old male

It makes me feel confident in our future.

– 53-year-old female

It makes me feel like I'm lucky to live in this lifetime.

– 35-year-old male

It makes me feel scared because then AI technology would have the upper hand on us.

– 33-year-old female

Makes me feel excited.

– 53-year-old male

Society

We build the machines, and we can control the limits of the AI.

– 37-year-old male tech exec

I see the value to it, but I also know with increased technology, there is always some negative effects.

– 43-year-old female, general population

AI could easily get out of control and overtake human control.

– 34-year-old male tech exec

I can see both benefits and potential problems with it. It all depends on how it is used.

– 55-year-old female tech exec

Pace of Technology Change Too Fast

Technology is changing our world at an astonishing pace, stretching our abilities to keep up. When asked to rate the pace of technological change, a plurality in both groups selected "too fast" and with similar percentages. In and of itself, that is interesting, but the strong correlation between these groups is noteworthy as even the tech execs feel the pace is too fast.

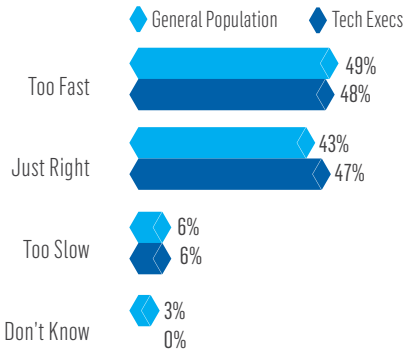

> *When can we just be happy with the way our current world functions now?*
>
> – 41-year-old female, general population

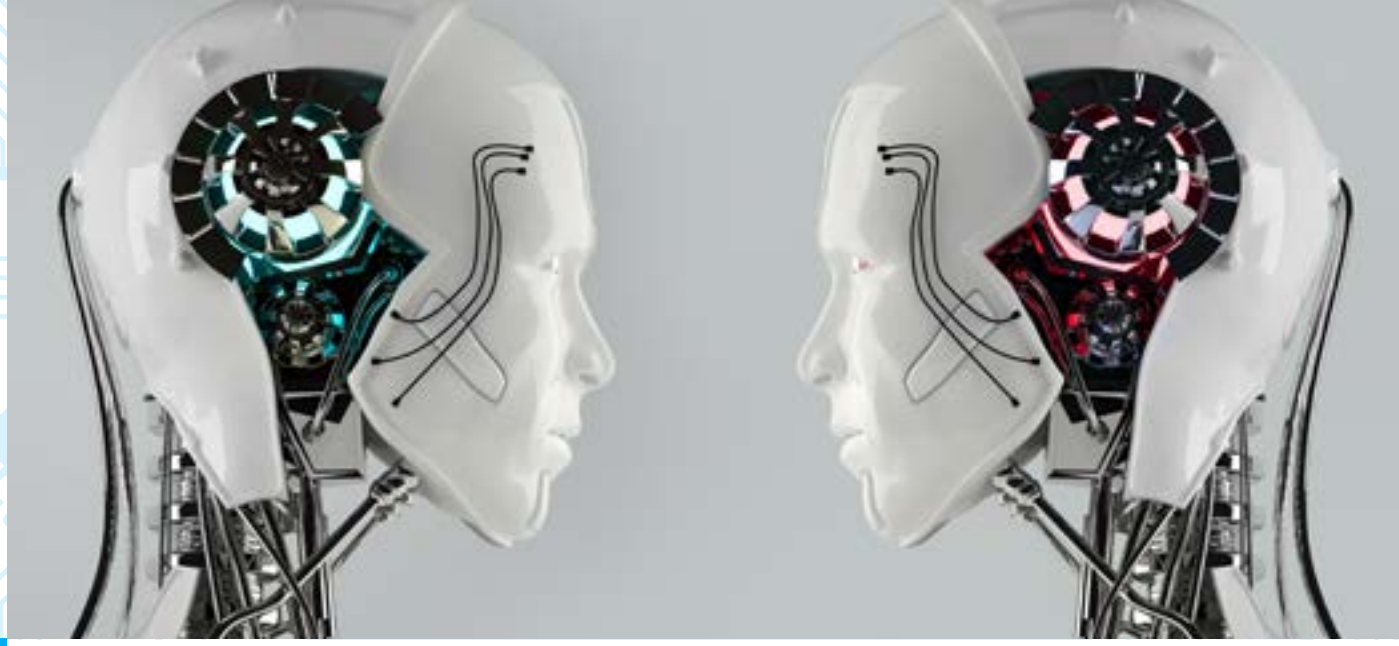

Two-Edged Sword of AI

A recent Dark Reading article noted how AI is revolutionizing cybersecurity for both defenders and attackers as hackers, armed with the same weaponized technology, create a seemingly never-ending arms race. This dual phenomenon of benefit and risk is evident in many venues where AI has found application. AI-powered virtual assistants such as Apple Siri provide increasingly personalized responses to our queries. Yet to achieve this they must collect ever more information about our interests and preferences leading to concerns about privacy.

AI is also useful for a new generation of digital video enhancement tools. One strain of these technologies enables a voice to be overlaid on an image and synced to realistic lip movement. A practical and useful application is to replace dubbing in films with a voice in the primary language of the viewing audience. Another is the de-aging of actors in film, to show the person as a young man or woman.

However, the same technology can be applied to create entirely fake audio and video. As shown in a Gizmodo story, these are becoming quite good and convincing. A Weekly Standard article notes these could be used to sow discord, undermine democracy and influence elections. As described in The Atlantic, the threat is that "manipulated video will ultimately destroy faith in our strongest remaining tether to the idea of common reality." In our survey, roughly a third of respondents believe fake videos could lead to war. About half believe these fakes could mean that no information is believable, corroding public trust in what they see and hear.

Of course, this double-edge is not new or unique to AI. Every wave of innovation creates winners and losers and brings aspects that harm segments of the population. For example, the advent of the automobile extinguished scores of businesses and professions, from carriage makers to teamsters. More recently, the internet has largely disintermediated the publishing industry especially impacting print newspapers and magazines. And online shopping is clearly having an adverse effect on traditional brick-and-mortar retail businesses.

The dichotomy between potential AI benefits versus possible harmful effects is evident in the Edelman AI survey results. Majorities of both survey groups, especially the tech execs, believe the positive from AI will outweigh the negative. Yet, significant concerns are also evident, such as worries the technologies will lead to greater social isolation. It is widely predicted that AI will lead to profound impacts just within the next five years, rapidly transforming the world such that society will struggle to address the many social, ethical and economic consequences.

Delta of Curiosity and Optimism Between General Population and Tech Executives

There have been many scary headlines about the potential impact of AI on everything from job loss and income inequality to an existential threat to humanity. Nevertheless, when it comes to the current state of the technology and the progress being made in AI, tech execs say they are excited (54 percent), curious (50 percent), optimistic (49 percent) and inspired (39 percent), while the general population say they are mostly curious (46 percent). This is in contrast to only 30 percent of the general population feeling excited and optimistic about AI.

HOW DOES PROGRESS IN AI MAKE YOU FEEL?

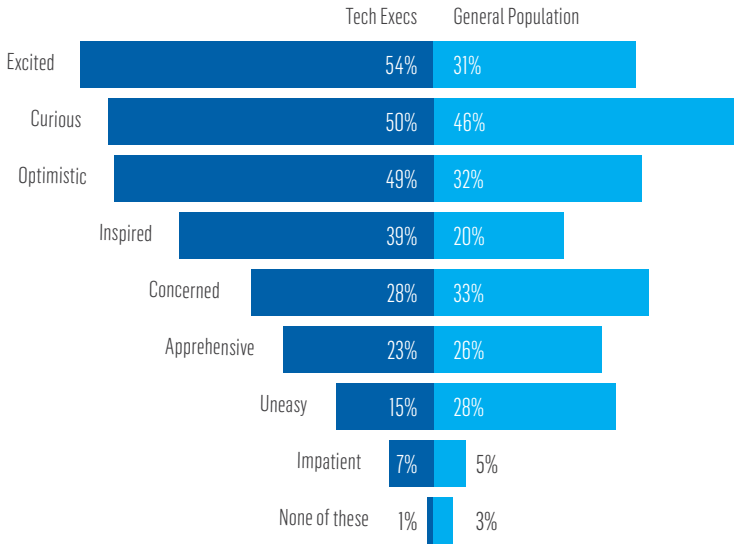

AI Already Impacting Work

AI is here now. Most respondents said AI is already having at least some impact on the work they perform. For tech execs, more than 8 in 10 said AI is already having an impact while 6 in 10 of the general population agreed.

"More robots are doing manufacturing jobs with human job losses.

– 39-year-old female, general population

DOES AI IMPACT YOUR WORK?

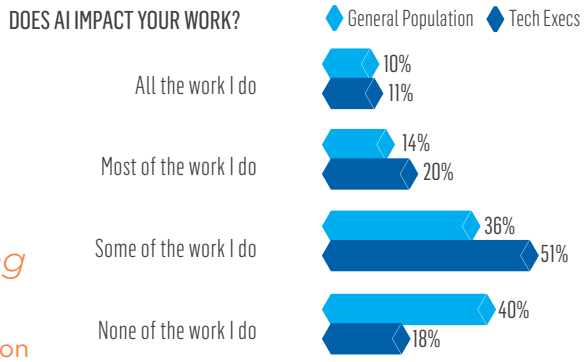

AI Work Impacts Grow within 5-10 Years

When looking out 5-10 years, 87 percent of tech execs and 68 percent of the general population believes at least some of their work will be impacted by AI.

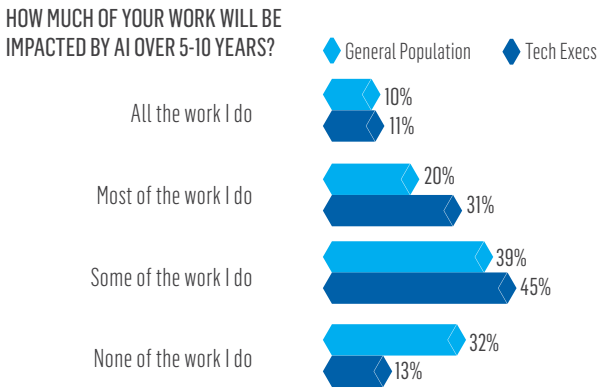

Groups Split on AI Job Impact

The potential impact of AI on jobs has prompted intense debate in industry and society and this is reflected in the survey results. Two-thirds of tech execs believe AI could increase the number of jobs. This opinion is not as broadly shared by the general population as only 45 percent of them hold this optimistic view. It is important to note that both groups agree that AI will cause job loss in the short term but will provide more jobs in the long term.

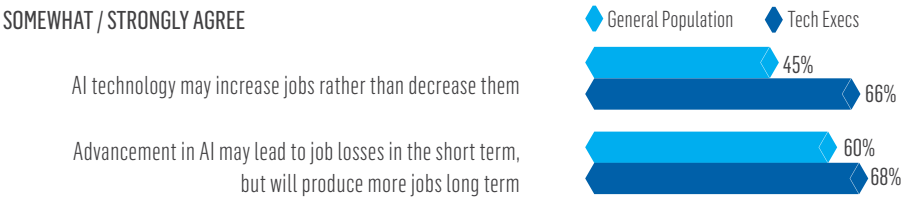

"The next big thing will be automatons and androids replacing human employees.

– 46-year-old female tech exec

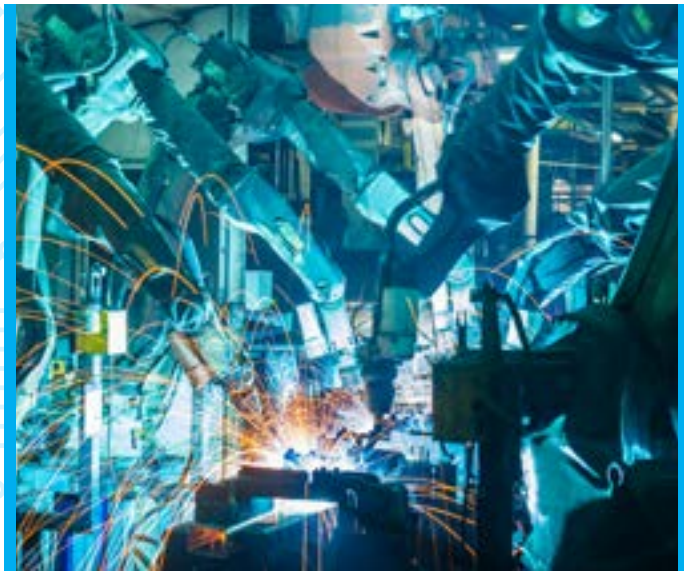

AI Will Benefit the Rich and Hurt the Poor

Both groups believe that AI will be beneficial for the rich, and there will be less benefits for others.

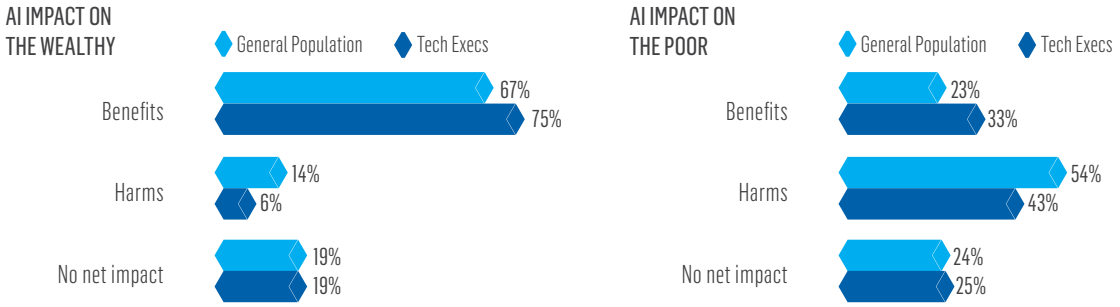

> *I'm nervous about the loss of jobs for the poor and middle class.*
>
> – 52-year-old male, general population

AI and Jobs

Perhaps no topic has been more speculated about than the impact of AI on the future of jobs, work itself and the potential for greater income inequality. Nobel laureate economist Wassily Leontief warned in 1983 that with the introduction of increasingly sophisticated computers, "the role of humans as the most important factor of production is bound to diminish in the same way that the role of horses in agricultural production was first diminished and then eliminated by the introduction of tractors."

Similar arguments have appeared more recently regarding the job-destroying potential of AI. A Pew Research study of "experts" found that 48 percent "envision a future in which robots and digital agents have displaced significant numbers of both blue- and white-collar workers—with many expressing concern that this will lead to vast increases in income inequality, masses of people who are effectively unemployable and breakdowns in the social order." These results are in-line with the responses of the general population in the Edelman survey where only 45 percent believe AI will increase the number of available jobs.

A Brookings Institute report cites Pew and several other studies to conclude 38 percent of current jobs could be automated. Their perspective is that due to job losses "Western democracies likely could resort to authoritarianism as happened in some countries during the Great Depression of the 1930s in order to keep their restive populations in check." Kai-Fu Lee predicts that AI technologies are going to displace about 40 percent of the jobs in the world. And unpredictable consequences from these shifts could well emerge.

Tech execs in the Edelman AI Survey have a much more sanguine view with two-thirds of respondents saying AI is more likely to increase the number of jobs. The optimistic argument is that technology revolutions always result in job displacement but ultimately create new industries and opportunities, and those industries will have to hire people. That has been the story so far, and it's been a positive story.

This view is consistent with a new World Economic Forum report that claims a net positive outlook for jobs between now and 2022, noting that "75 million jobs may be displaced by a shift in the division of labor between humans and machines, while 133 million new roles may emerge that are more adapted to the new division of labor between humans, machines and algorithms." Similarly, a PwC study asserts that AI, robotics and related technologies should "generate enough new jobs to broadly offset the potential job losses associated with automation." This view is echoed by analyst firm Gartner, which stated that AI-related job creation will reach a net 2 million new jobs by 2025.

However, the time between now and 2025 is only the first phase of what will be a dramatic shift. A further study from PwC describes three overlapping cycles of automation that will stretch into the 2030s, each with their own degree of job impact. In a further report from McKinsey, the firm notes that automation and AI will lift productivity and economic growth, but millions of people worldwide may need to switch occupations or upgrade skills.

Going forward, there is really no certainty about the jobs outlook in the AI era where machines will increasingly do the work for us. In the Edelman Survey, more than 7 in 10 respondents in both groups said there were too many variables to know the net effect of AI. What is clear is that the transition to the AI economy will be uneven and many people will be left behind. In response to these trends, futurist Martin Ford speculated in a TedTalk that we could very well end up in a future with significant unemployment and "soaring levels of inequality." To help offset this, he and others believe it is time to seriously consider providing people with a guaranteed income or universal basic income.

Fears of Social Isolation, Digital Group Think

More than 70 percent of both groups agree that AI will lead to greater social isolation – less in-person interaction – and nearly the same percentage believes it will lead to a loss of human intellect and of creativity, with the potential for digital group think.

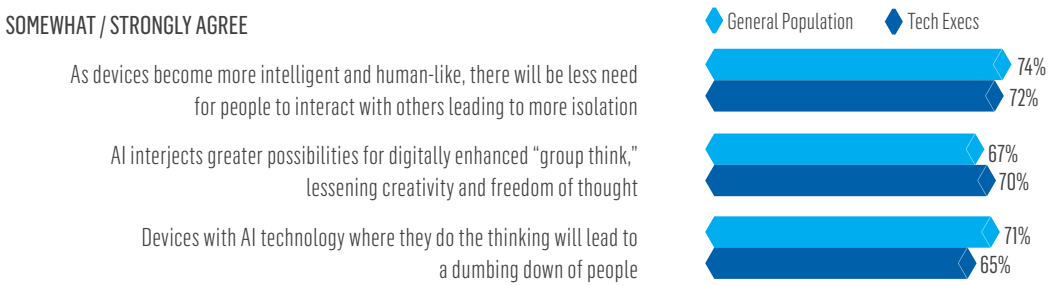

> *AI is something made to be helpful but [instead] will make people lazy.*
>
> – 27-year-old male, general population

Human Intellect Could Atrophy

Fully 9 in 10 of the tech execs somewhat or strongly agree that AI-powered machines will take on mundane tasks allowing people to have greater freedom to pursue more creative work and play. More than 3 in 4 of the general population group agrees.

If AI will do so much for us, from better medical diagnoses to finding criminals, might this lead to a dumbing-down of humanity, characterized by less curiosity and competence? More than 7 in 10 of the general population group thinks so, as do 65 percent of tech execs. The concern is that we will learn to lean on our smart devices and applications for so many things, we could become less inquisitive and more trusting of the information we are provided as accurate and authoritative.

A worry is that people will move through life on autopilot, just like our cars, in the not-too-distant future. If smartphone use is any indication, there's some cause for worry about AI. Nicolas Carr wrote in the Wall Street Journal about research suggesting the intellect weakens as the brain grows dependent on phone technology. Likely the same could be said for any information technology where content flows our way without us having to work to learn or discover on our own. If that's true, then AI, which increasingly presents content tailored to our specific interests, could create a self-reinforcing syndrome that weakens our intellect. This concern is echoed by Daniel Weld, a professor at the University of Washington who studies human-computer interaction, in an Axios article: "I worry that human abilities may atrophy."

Part of human drive and creativity is to test ourselves against our environment, notably other people. Recently, AI has proven to be more proficient than people in many areas. When people cannot win, it is possible they will simply stop trying. The positive side to this story is that people may learn to give unto AI what it can do better and to instead focus on what we can do better. This fits with Moravec's paradox, which states that tasks we find difficult are easy for a computer, and vice-versa. Or, as Kai-Fu Lee says in a Forbes article, "Let us choose to let machines be machines, and let humans be humans."

AI-Powered "Deepfakes" Could Lead to War

People are concerned about the impact of deepfakes," where AI-tools are used to falsify reality and manipulate perception through faked audio and video, potentially leading to reality apathy, human puppets and the "Infocalypse." When asked about the impact, about half of both groups believe deepfakes will further erode public trust and 1 in 3 believe deepfakes could lead to a shooting war. Determining authenticity of content will be increasingly critical.

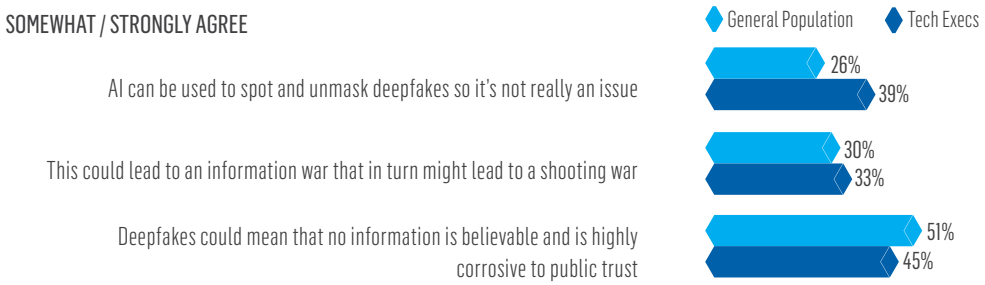

Deepfakes Corrode Trust

Deepfakes are videos or audio recordings that are doctored to alter reality, "showing" events or depicting speech that never happened. Developers are using deep learning technology - and thus the term deepfake - to identify the facial movements of a targeted person and then render highly realistic, computer-generated fake videos with real-looking lip movement, facial expression and background and audio. The most recent techniques produce fake videos that are nearly indistinguishable in quality from the source materials. These videos are still not perfect but improving rapidly and will soon be very convincing even to an expert eye and ear.

These fakes mean it is now possible to portray someone – say the leader of a country - saying pretty much anything the video creator wants and in whatever setting they want. A recent example is an altered video of President Trump that was broadcast by a Seattle television station.

The implications of deepfakes are staggering. The line between what is real and not real is changing. Many of the early deepfakes are pornographic, with developers replacing the faces of porn stars with those of celebrities or even the person next door. Beyond that, fake videos could falsely depict an innocent person participating in a criminal activity, falsely show soldiers committing atrocities or world leaders declaring war on another country and possibly triggering a very real military response. Because people tend to lend substantial credence to what they see and hear, deepfakes could soon become a very real danger. Just as with "fake news," if fake videos and audio become common on social platforms and websites with extreme agendas, people may start questioning real videos. At that point, people may lose the ability to discern what is real and what is not, what is truth and what is fake, and reality will lose its meaning. This leads to any number of concerns, not the least of which is the question of whether people are capable of living in a world where there is no credible "truth." This could easily undercut our basis for rational decision making.

Researchers are at work developing approaches also using AI that could identify these fakes. Recent advances pinpoint eye blinking as a weakness in fake video development. While software could soon be available to detect this weakness, it is very likely that deepfake developers will then improve their technique. Much like cybersecurity, where hackers and those trying to thwart them leapfrog one another, the same will be true for deepfakes. For now, at least, people will be on their own to discern fact from fiction.

Downsides to AI-Powered Smart Toys

While both groups see AI being helpful for developing smart toys, they view the impact with a mixture of approval and worry. Regarding the latter, fully 7 in 10 believe these toys could take the place of human friends and would invade children's privacy.

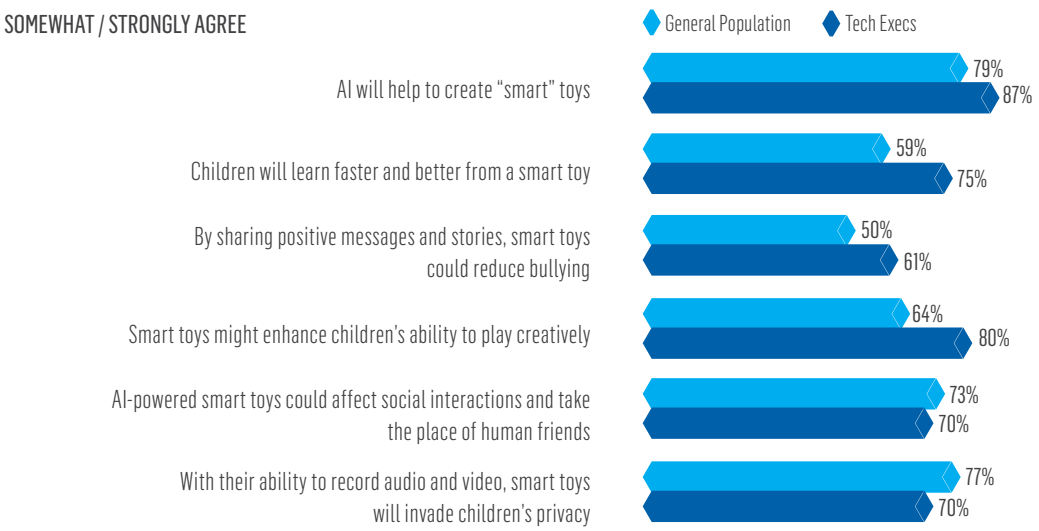

Smart Toys Not That Simple

Electronics in toys are far from new, having arguably begun with the Little Professor from Texas Instruments in 1976. However, AI technologies are now being incorporated in the new generation of toys. These smarter toys interact with their environment including people and objects, with the toys adjusting their behavior as the software learns. Through use of machine learning, the toys can "understand" what is happening around them quickly, making it more interesting for a child to play with. Today these toys range from robots to race cars, from dolls to drones and new AI-enabled toy products are continually coming to market.

While these toys may be fun, there are concerns being raised about unintended consequences. For example, there's a doll that uses AI-powered speech recognition technology to converse with a child. This doll records the conversations and transmits these online to a voice recognition and analysis company. These conversations in turn are fed into a deep learning algorithm leading to improved conversational abilities for the doll. However, privacy issues have led regulators in at least one country to remove the toy from their market. The doll is not the only such toy with these voice capabilities.

The World Economic Forum has identified several issues with AI toys including privacy, bias, surveillance, manipulation, transparency and accountability. Questions are also being asked about the longer-term effect of playing with smart toys, both cognitive and social. To the latter point, some are asking about the possible implications of a toy becoming someone's best friend. After all, these are artificially intelligent products programmed to produce affiliation and affection. A concern is that it will be easier or more fun to interact with a toy than a real person, thus a child may spend less time interacting with other children. The counter to this is that playing with a smart toy may not be much different than when a child has an imaginary friend. However, most children outgrow the imaginary friend, but it might be different with smart toys.

Supporters of smart toys highlight the educational benefits and claim they could be especially useful to close the learning gap for the developing world and for those with learning disabilities. They argue the toys will enhance learning by helping young children recognize and resolve conflicts, teaching those with special needs the social skills to successfully interact with others, and assist with math, grammar and language skills.

AI Could Help the Elderly

Survey respondents are generally positive on how they see the impact of AI on assistive technologies for the elderly. Specifically, AI could help them live independently at home, but there is also a belief that AI will be too complicated for the elderly to use. There is modest support for the idea that robots will be needed as a replacement for human caregivers.

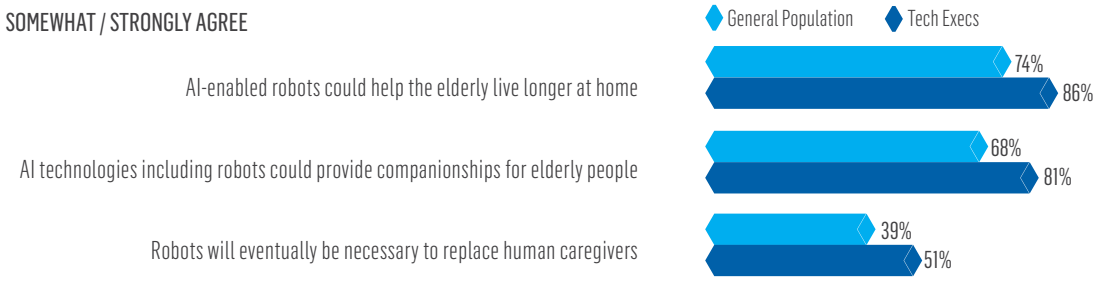

> *Hopeful that these inventions will provide needed help for the elderly.*
>
> – 37-year-old female, general population

AI Robots to Help Aging Population

The number of Americans age 85 and older is projected to more than double in the next few decades, soaring from 1.8 percent of the population in 2010 to 4.5 percent in 2050. Yet the demand for professional caregivers already far outstrips supply. The aging population, both in the U.S. and in many other developed countries, will increase demand for elder care assistance. AI, in large part through the application of wearables and robotics, could help with caring for elders, increasing their independence and possibly reducing the detrimental impact of social isolation.

Dreams of a robotic workforce are not new, arguably beginning in 1951 with Nobel Prize-winning physicist William Shockley when he decided to pursue development of an "automatic trainable robot." Nearly 70 years later, that vision is coming to fruition. In the foreseeable future, robots will be able to help with myriad daily activities, from assistance with opening bottle lids to mowing lawns, folding towels, lifting people out of a chair or bed, fetching food, turning lights on and off, and mopping floors. Much of this is already available and in use, especially in Japan where individuals over the age of 65 make up more than a quarter of the total population, and there is an acute shortage of human caregivers. In Canada, a robot is now being tested in nursing homes by engaging in conversation with residents to determine signs of dementia.

More than 8 million Americans over 50 are already affected by isolation. Loneliness is a significant predictor of poor health, and it is widespread, affecting more than one-third of older adults in the U.S., according to a 2010 AARP study. This isolation may contribute to many diseases and impair immune systems. Researchers have suggested the risk of dying from social isolation is on par with smoking 15 cigarettes every day and being an alcoholic. Northeastern University professor Timothy Bickmore says in The Wall Street Journal, "Robots that help people connect with and maintain their relationships with others are becoming increasingly important." While early evidence suggests that robots can provide effective social interaction, they are not thought to be a suitable replacement for human touch and care. Even with further advances in robotics, they will likely fall short of human-like qualities and unable to give satisfying emotional support.

Concerns also exist about the impact these eldercare robots may have on jobs for those working in the industry. If the economics of robots improve to where they cost less than human caregivers, there will be job displacement. As most of the caregivers are women, this group could be disproportionately affected by the growing use of robots for eldercare.


Business & Government

Just hope we use our AI powers for goodness and not evil as the possibilities for serious abuse are there.

– 58-year-old male, general population

AI is inevitable, looking at how technology is advancing, but I always believe in the human factor.

– 48-year-old female tech exec

High Stakes

AI is evolving at an extraordinary rate. As Elon Musk pointed-out, AI is advancing fast due to the double exponential growth in both computing power and software capability. "... the degrees of freedom to which artificial intelligence is able to apply itself is really increasing I think by 10 orders of magnitude a year." With the coming of advanced hardware designed specifically to accelerate deep learning algorithms, that rate could soon get another dramatic boost.

Many technology companies see the potential for AI and have moved to develop machine learning and deep learning applications. These, in turn, are being adopted by companies across many industries as business sees the potential benefits for greater insights, efficiency and servicing their customers. Startup companies looking to capitalize on the AI trends have proliferated over the last few years. Taken together, the drive toward AI resembles a gold rush.

Amid this stampede, some are going to be trampled. Hence there have been dire warnings about job losses. While the absolute number of jobs impacted is an open question, it is widely agreed that millions of people will need retraining or will have to find new professions. That could be relatively easy for some of those displaced but much harder for others, and there is the possibility that AI advances could leave many behind and create a new permanent underclass. In response, ideas about guaranteed minimum incomes are being widely discussed, with the financial source possibly coming from a tax on AI development and applications.

Beyond jobs, there are myriad issues surfacing about AI ranging from inherent algorithm bias to the need for greater transparency about how the technology works. In the United States, current government policies at the federal level amount to a laissez-faire approach, leaving it largely to private companies and the marketplace to sort through these challenges. Many companies are participating in these conversations for socially responsible use of AI through consortiums such as the Partnership on AI and the Information Technology Industry Council. Several well-known companies have also recently outlined their AI code of ethics with some even asking for more government regulation. Musk stated that the right move now is to establish a government regulatory agency to ensure public safety, with a role like the FAA and FDA.

Of course, AI has become a global horse race with multiple countries jockeying for leadership. The United States has been in a position of strength due to its academic institutions and breadth of technology leaders within its borders. China, however, is making a concerted effort to become the world leader in AI. They are not alone. For example, a Quartz article claims AI is the new space race and describes the efforts of France, the U.K., Canada and Russia to advance their AI capabilities. India too, is making AI a national priority. The U.S. government only recently has made moves in this direction with the formation of a National Security Commission on Artificial Intelligence that includes 15 technology experts headed by former Alphabet Chief Executive Eric Schmidt. Among the group's areas of interest are research funding, workforce reskilling and AI ethics. Only last month, President Trump issued an executive order directing greater support for AI efforts. Though very light on specifics or funding, the order covers technological development and progress in related areas including government regulation, education, private-public sector data sharing, and international diplomacy. To further highlight the importance of AI, Russian president Vladimir Putin said in a 2017 speech: "Artificial intelligence is the future, not only for Russia, but for all of humankind." He concluded by declaring that whatever country comes to dominate this technology will be the "ruler of the world."

The stakes for government and business could not be higher.

AI Benefits All Industries

Majorities in both groups see positive benefits from AI across all industries. In addition to the technology industry, respondents see the most positive impacts for manufacturing, energy, telecommunications, biotech / life sciences and healthcare. Interestingly, financial services rated lowest by both groups, yet there are many AI applications already being deployed from stock selection to credit ratings.

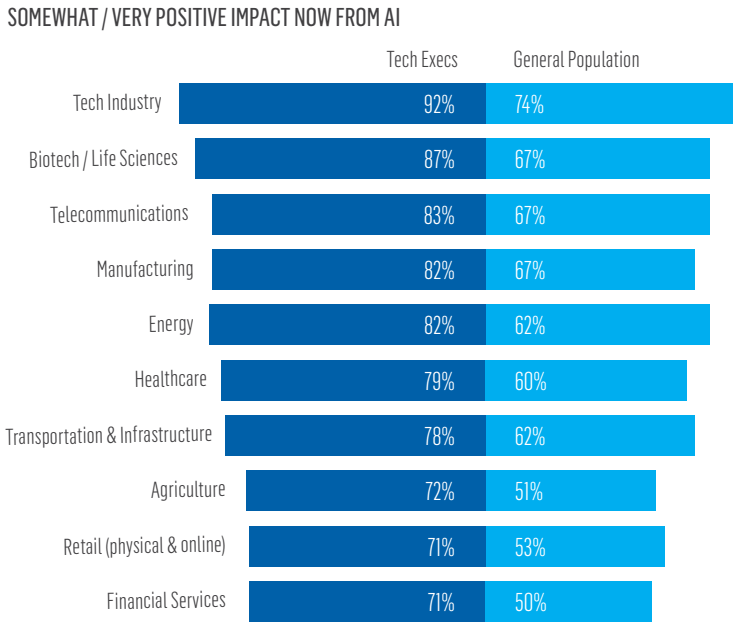

> *Hopefully AI can develop effective treatments for genetic diseases.*
>
> – 22-year-old female, general population

Technology Industry Is Under-Regulated

Support for regulation of the technology industry extends beyond AI concerns. Perhaps surprisingly, even tech execs acknowledge that the industry has been under-regulated.

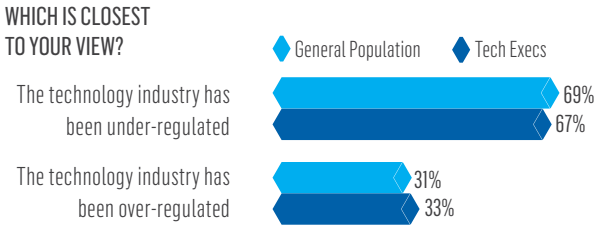

Majorities Say AI Regulation Is Critical

At present, there is no real regulation for AI development or use. More than half in both groups responded that AI regulation is critical. Most believe this should be managed by a public oversight body, while few believe that those developing and deploying AI should regulate themselves.

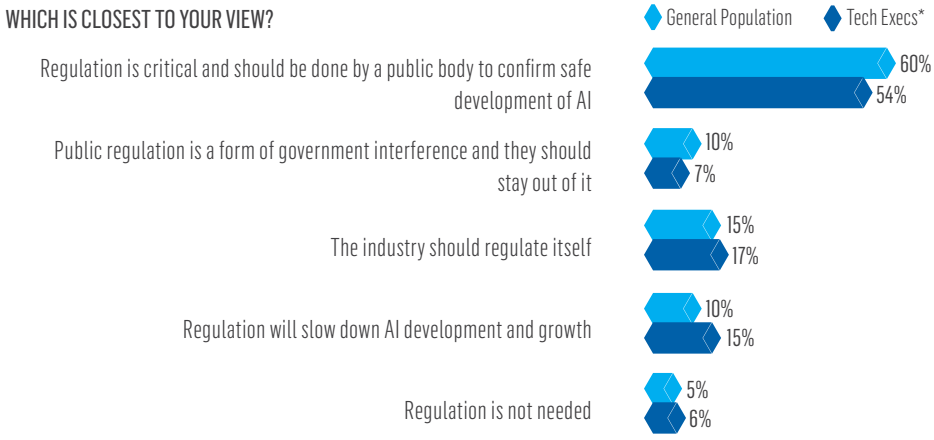

* Percentages do not always add to 100% due to rounding.

Need for AI Oversight

Both groups strongly believe human oversight of AI is a much needed requirement. Only half of respondents believe the federal government should have an active role in determining and influencing policies for managing AI-powered applications. A strong majority of both survey groups believe the tech industry should be active in this capacity.

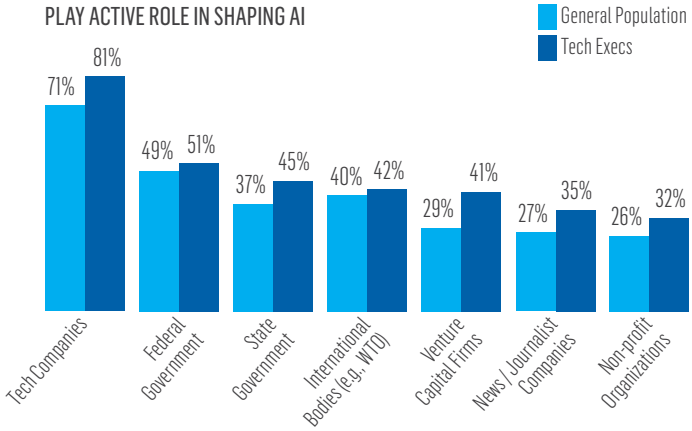

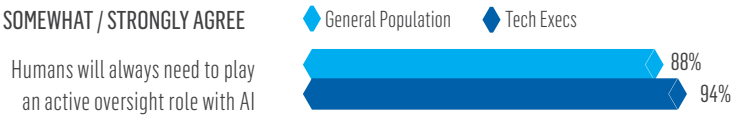

Is Now the Time for AI Regulation?

Majorities in both groups (general population: 60 percent, tech execs: 54 percent) responded that regulation of AI is critical and needs to be done by a public body that has insight and oversight to confirm safe development of AI. Per a GeekWire article: "This isn't as simple as just 'trust,'" according to Kay Firth-Butterfield, Project Head for AI and Machine Learning at the World Economic Forum's Center for the Fourth Industrial Revolution. "This is more complex, because the technology itself is very fast, changing all the time and is complex as well."

There are any number of areas where regulation could help. For example, loan analyses including credit card applications are now often performed using AI algorithms. Yet, how can an algorithm be held accountable if a customer feels that a decision about their credit card application was wrong? Many argue that people have a right to know how decisions that affect them are being made. Care also needs to be taken to provide transparency into how the algorithms work to ensure there are no inherent biases. Similar needs exist across many potential applications ranging from AI use in protecting critical infrastructure to enabling autonomous vehicles and protecting privacy rights in healthcare.

For its part, the technology industry is becoming more comfortable with AI regulation. For example, Microsoft recently urged lawmakers to regulate the use of AI-powered facial recognition software to prevent abuse. In a thoughtful blog, the company makes the point that "advanced technology no longer stands apart from society; it is becoming deeply infused in our personal and professional lives" and called for government regulation aided by bipartisan, expert commissions to prevent abuse. Similarly, Google has released a set of guiding AI ethics principles.

The challenge for regulation is to balance potentially competing factors of technology development and consumer protection. Government and the industry may be able to provide a stable and transparent regulatory environment to develop AI, while overcoming or at least assuaging the potential fears about these developments. Given the rapid pace of AI development and the potential impacts, combined with the critical requirement according to survey respondents, there is a need to move quickly to put common sense regulations into play.

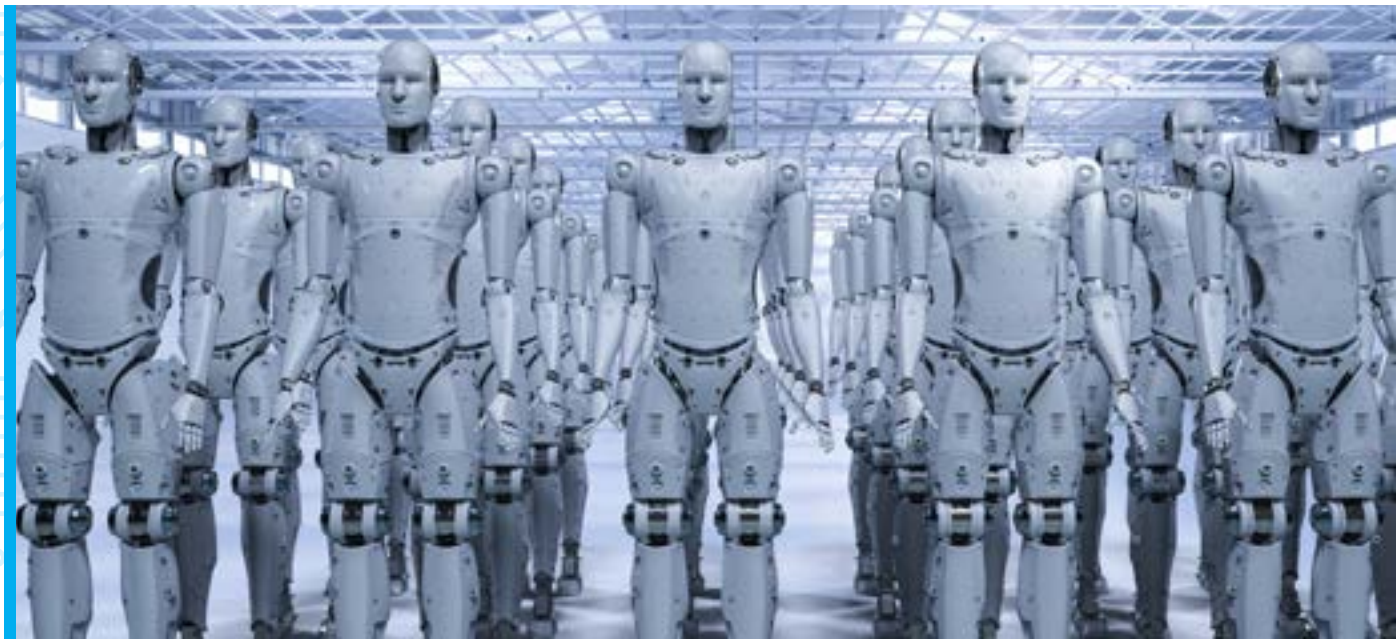

Technology Companies Must Practice Responsible AI

More than three in four respondents in both groups are concerned that technology companies aren't thinking through the long-term consequences of AI development, contributing to bubbling concerns around AI technologies.

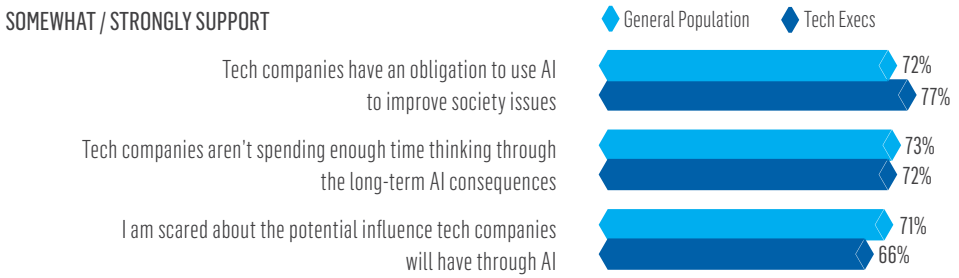

Societal Value Required for AI Businesses

Respondents in both groups trust the business community will prioritize AI profits but also see important roles for companies including innovation for societal benefit and skills training. Companies have a tacit license to operate and to make money from AI, but this must be accompanied by creating societal value.

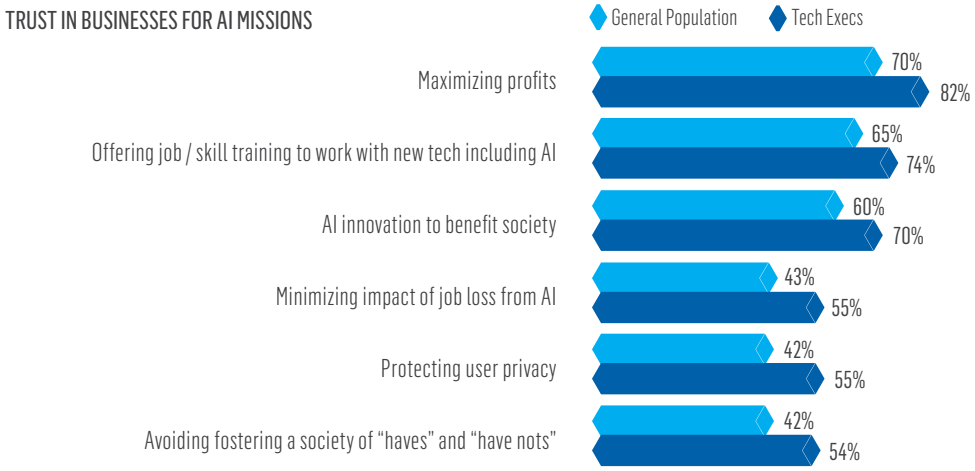

> *AI needs to help us solve today's issues that threaten our lives.*
>
> – 29-year-old male, general population

Addressing AI Concerns

Given the uncertainty about AI development and concerns about the impact, a majority of both groups believe that transparency in how it works is important along with demonstrating how the technology will improve lives and create jobs. Only about a third of respondents felt it is important to diversify the AI development workforce.

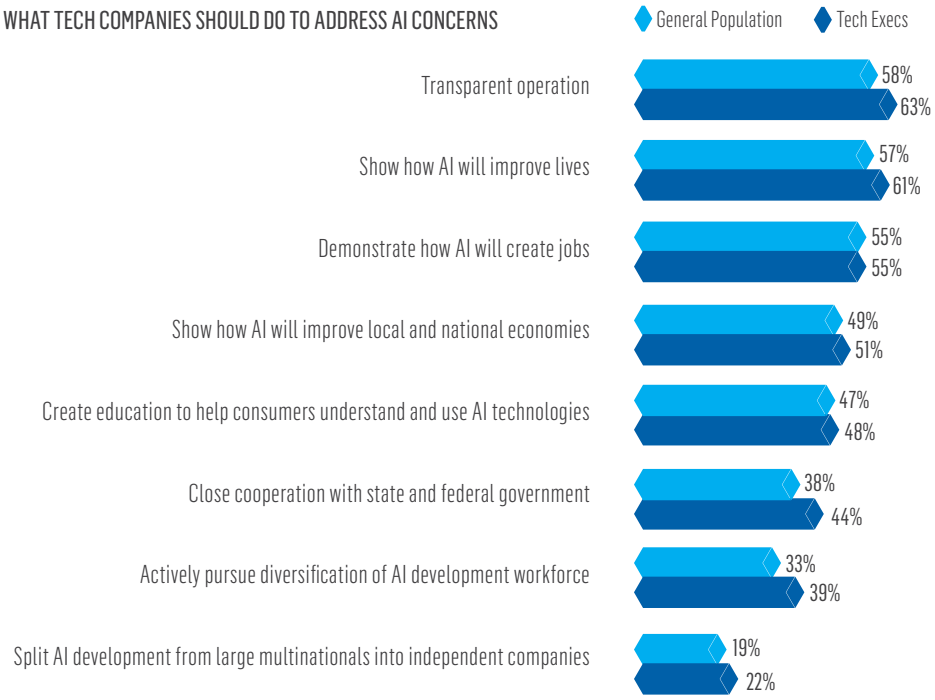

Responsible AI Development

One of the drivers behind the growing demand for regulation of AI is a perceived lack of transparency into how the technology works. Some AI algorithms are referred to as a "black box," meaning their operation is opaque to the outside user. It turns out that these could be black boxes even to those who developed the algorithms. This has led to a few voices charging that AI today is little more than alchemy – a seemingly magical process.

Others have claimed that certain AI algorithms incorporate inherent (though perhaps unconscious) bias. A prominent court case is centered on a man who was sentenced to prison in part because of content produced by a "secret algorithm." According to New York Times reporting, Wisconsin's Attorney General said the defendant was free to question the assessment and explain its possible flaws. However, the algorithm is proprietary and – even assuming its operation could be explained – is not available.

Majorities of survey respondents said that technology companies need to provide transparency with respect to AI operation. The industry is beginning to move in this direction. For example, IBM Research has published their strategy to address multiple dimensions of trust to build AI solutions that inspire confidence. According to an article in Futurism, the company has proposed that AI algorithm developers create a document that describes how well the algorithm performed in standardized tests of performance, fairness and risk factors, and safety measures.

These moves towards greater transparency are laudable but may not ultimately be entirely practical. As noted by Wired, AI algorithms often use deep-learning techniques to analyze data with thousands of variables, arranging them into immensely complex arrays of relationships and then running those arrays repeatedly through neural networks. These then produce the often-astonishing outcomes such as better medical diagnoses. The article notes that our brains simply could not contain this much information and implies that we might not ever fully understand how the algorithms work.

Which is not to say that seeking transparency is a fool's errand. But it does suggest, however, that the quest for responsible AI development is heavily nuanced and will require close cooperation between the various stakeholders. Members of the public, government and business need to engage in an open and honest debate about what needs to be made transparent and how best to achieve this within the constraints inherent in the complexity.

Six Steps for Businesses to Earn AI Trust

AI is transforming business and industry and will lead to substantial economic growth. According to a McKinsey Global Institute report, AI technologies will deliver additional global economic activity of approximately $13 trillion by 2030 – or about 16 percent higher cumulative GDP compared with 2018. However, AI will also have profound impacts on the social fabric. Some number of job categories will become obsolete, and new job categories will be created, though for many years the results will be uneven with periods of time where job losses could be extensive. A new World Economic Forum report claims a net positive outlook for jobs between now and 2022, noting that "75 million jobs may be displaced by a shift in the division of labor between humans and machines, while 133 million new roles may emerge that are more adapted to the new division of labor between humans, machines and algorithms."

However, a PwC report describes three overlapping cycles of automation that will stretch into the 2030s: the algorithm wave, the augmentation wave and the autonomy wave. PwC notes that in the first algorithm wave, projected to extend into the early 2020s, automation will replace relatively few jobs. The more impactful waves are expected to come in succession from the mid- to late-2020s and continue into the mid-2030s. According to PwC's findings, automation will impact 30 percent of jobs during this timeframe. Other estimates are for even higher job displacement. The political and economic consequences of this are potentially dramatic.

AI offers great promise, but for many, it will also bring painful dislocation. During times of dramatic change when emotions will run hot, business could become a target, not only for basic regulation but potentially for antitrust investigations and consumer actions such as boycotts. Beyond the clear need to demonstrate algorithm transparency, it is incumbent on businesses to act with a social conscious. This is always good advice, but perhaps never more so than now for businesses developing or deploying AI technologies. Following are six key actions companies can take to enhance their leadership and protect their reputation during times of profound change.

Be proactive with policy-making. Companies can wait for events to overtake them and lead to less than desirable outcomes, or they can actively participate in critical policy decisions. A clear result from the survey is a call for AI regulation. Microsoft took an unusual step for business by agreeing that at least in some instances, AI regulation would be useful. In a blog post from company President Brad Smith, Microsoft called for both corporate responsibility and public regulation for facial recognition technology, which so far has proven to be more accurate for white men than for women or people of color. By establishing this point-of-view, the company demonstrated leadership on an issue of critical importance and has enhanced their public credibility.

Establish and adhere to an AI ethical code. AI use will only expand in the future, inevitably causing many ethical issues to arise as algorithms increasingly operate cars, homes and businesses. Businesses should create a code of ethics regarding the use of AI technologies to ensure their behavior is above reproach and to be better stewards of public trust. IBM, for one, has started doing this with their "Everyday Ethics for Artificial Intelligence." Similarly, German software powerhouse SAP recently released an ethics code to govern its AI research, aiming to preserve integrity and trust by preventing the technology infringing on people's rights, displacing workers or inheriting biases from its human designers.

Perform and document rigorous algorithm testing and provide transparent operation. Machine learning algorithms are complex mathematical formulas and procedures and have an increasing impact on people's lives. As decisions become increasingly governed by these algorithms, the decisions they influence are increasingly opaque and less accountable. It is incumbent on business to be as transparent as possible to explain the operation of their algorithms.

Demonstrate responsible actions to minimize negative impacts from AI. Business leaders will increasingly need to develop a thorough workforce plan to meet the challenges of this new AI-powered era. This includes a focus on providing skills training to survive and thrive in an AI world. Improving access to education for their employees – including tuition reimbursement and other incentives – is useful. Companies should develop a position on guaranteed minimum income programs, which are already being tested in several countries.

Showcase societal benefits of AI application plus real-world impact at scale. People like the benefits that AI currently offers such as speech and image recognition, search engines, spam filters, product and movie recommendations, and more. However, negative AI use cases also exist including security hacks, phishing scams, deepfakes and personalized disinformation campaigns. Most would agree that the use of AI technologies is fine if there are clear societal benefits. That could be relatively easy to demonstrate for advanced medical uses but likely not as readily evident for some other fields such as supply chain management. In this latter case, there are benefits to leveraging the vast amounts of data collected by industrial logistics, warehousing and transportation systems to increase efficiency, reduce costs and potentially lower prices. Businesses deploying AI would benefit by clearly articulating the value the technology offers to their employees and end customers.

Don't hide your light. It is often best to be proactive when taking the steps outlined above and communicating these perspectives with key stakeholders through blog posts, bylined articles, feature stories and public speeches. In doing so, companies are on record as acting responsibly on behalf of their constituencies and society. This goes a long way towards earning trust and serves as an opinion buffer should something go wrong at a future point in time.

Conclusion

The 2019 Edelman survey revealed that people are generally optimistic about AI, especially so for tech execs. Both they and the general population group believes AI has the potential to transform business, industry and some social services to an unprecedented degree and with astonishing speed. This augurs well for productivity and will be a boon to those who will directly benefit. Respondents also believe that artificial general intelligence – also known as superintelligence or the singularity – will arrive within 5-10 years, which is decades sooner than the views of most AI experts.

At the same time, the survey revealed worries the AI revolution could produce tumult with job losses that would hurt the poor and lead to societal disruptions. If autocratic powers are able to access and leverage big data content processed and personalized with AI technologies, they could increasingly control the levers of disinformation thereby manipulating and controlling their populations. About 7 in 10 believe AI interjects greater possibilities for this digitally enhanced "group think," lessening creativity and freedom of thought. Public trust of content is also threatened with the rise of AI-powered deepfake videos as evidenced by a third of respondents' concern that these could falsely lead countries into war.

In summary, the AI future is rich with possibility but also contains significant risks. As Darrell West of the Brookings Institute notes, AI has the "potential to move civilization forward in progressive ways. But without adequate safeguards or the incorporation of ethical considerations, the AI utopia can quickly turn into dystopia." It is clear from the results that business and government need to be thoughtful and take action to ensure the potentially harmful effects of AI do not cause society to lose the clear benefits of the technology. Reflecting majority views from both groups, this includes regulatory oversight of AI development and implementation. There was additional agreement among respondents that it is incumbent on technology companies producing AI products to demonstrate transparency with respect to how AI operates and to show how it provides positive societal value.

In sum, we discovered at least 10 noteworthy takeaways from the survey.

Key Takeaways

1. AI is creating a technology revolution with benefits across all industries.
2. Tech execs are far more optimistic about AI than is the general population.
3. Strong majorities believe that AI will benefit the wealthy while nearly half expect the poor will be harmed.
4. It is widely believed that AI will lead to short-term job losses with potential for societal disruption.
5. AI can be leveraged as a tool of division by autocrats, sowing disinformation to create digital group think.
6. AI-enabled "deepfakes" add to the "fake news" narrative and will fundamentally undermine trust in what is real and could lead to war.
7. Concerning impacts from AI deployment include more social isolation and less human intellectual capability.
8. Half of both groups believe the pace of technology is advancing too quickly.
9. Tech companies must demonstrate AI transparency and social benefits.
10. There is a clear need to do more to regulate AI development and deployment.

Appendix

Survey Methodology and Profile of Tech Executives

Methodology

The survey was conducted in the summer of 2018 with 1,000 U.S. adults from the general population and 300 executives working in technology roles. The questions in the survey were developed by the Edelman AI Center of Expertise and augmented with input from the World Economic Forum. The general population was equally distributed from around the country and selected irrespective of their work positions while selection of the technology executives focused on those who held senior management and C-Suite positions in their organizations.

TECH EXECS	
Senior Executive Level (CEO, CTO, CSO, President)	15%
Executive Level (Executive Vice President, General Manager)	10%
Upper Level Management (Vice President, Senior Vice President)	18%
Mid-Level Management (Director, Senior Manager)	50%
Company Owner / Sole Proprietor	6%
Part Owner (at least 25% equity)	1%

COMPANY SIZE	TECH EXECS
1 - 24	12%
25 - 49	3%
50 - 149	11%
150 - 249	11%
250 - 499	9%
500 - 999	15%
1000 - 1499	11%
1500+	29%

Tech Executive Roles

43 percent of the technology professionals surveyed held positions of vice president or higher in their firms, and 93 percent are senior managers and above.

55 percent of the tech execs are from companies with more than 500 employees.

Tech Executives Familiar with Leading Technology

Tech execs were far more likely to be familiar with leading edge technologies.

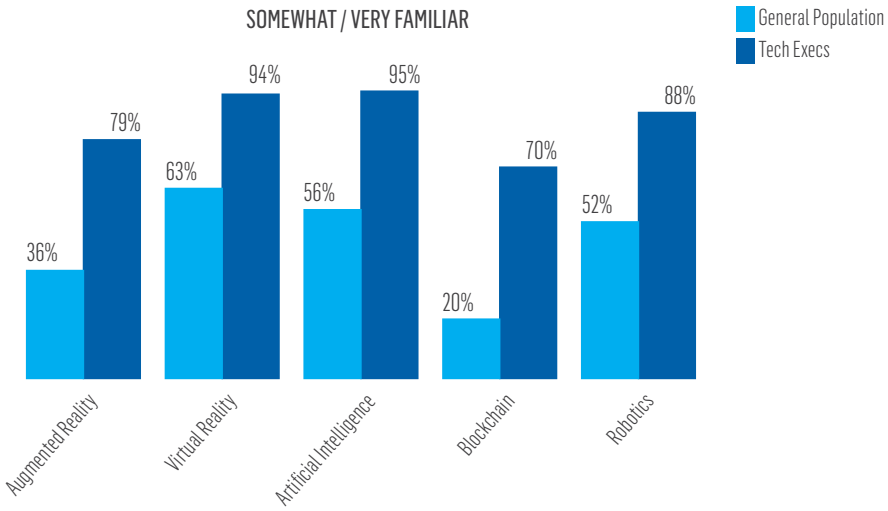

Tech Executives Are Early Adopters of Tech

81 percent of the tech execs group consider themselves early adopters of products based on new technology compared to only 30 percent of the general population group. Tech execs also reported they used more computing devices than did the public at large demonstrating they were far more comfortable with technology.

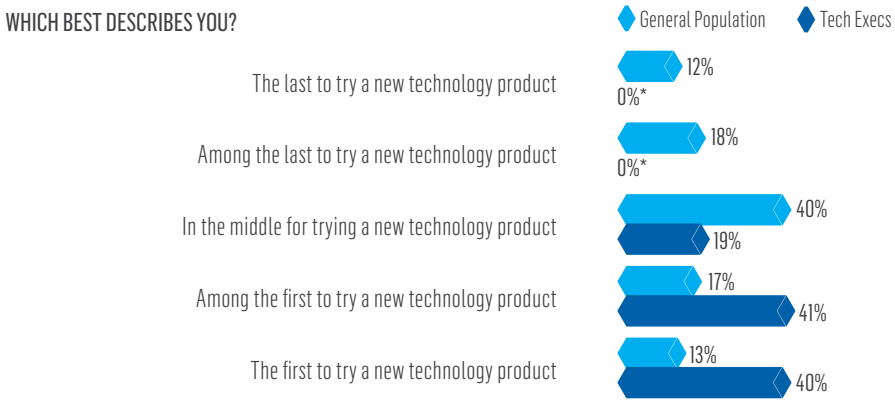

* Our focus in the survey for tech execs was on leading-edge technology adopters. Respondents were disqualified from participation if they were among the last to try new tech products.

Tech Executives Have More Education

By a significant margin, tech execs are better educated. Regardless of that difference, the philosophical perspective of the two groups are remarkably similar with strong pluralities claiming a moderate view.

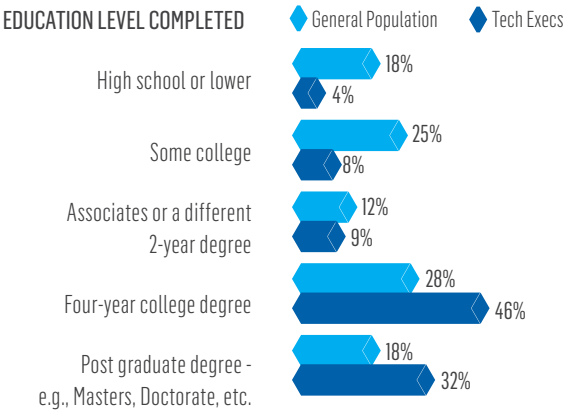

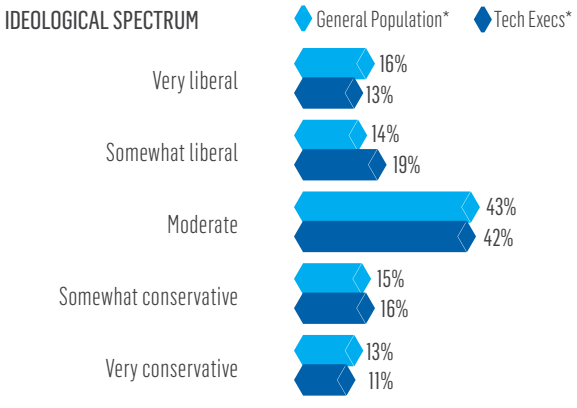

* Percentages do not always add to 100% due to rounding.

Tech Executives More Closely Follow News

These tech execs were considerably more likely to closely follow news on a variety of topics. Not surprisingly, this is especially true for technology and science issues but was consistent across a broad spectrum.