

2018 AI predictions

8 insights to shape
business strategy



pwc

pwc.com/us/AI2018

Here's some actionable advice on artificial intelligence (AI), that you can use today: If someone says they know exactly what AI will look like and do in 10 years, smile politely, then change the subject or walk away.

AI is remarkably complex and advancing quickly. It's doing far more in some areas, and far less in others, than anyone would have guessed a decade ago. It's impossible for anyone today to give a precise vision of how the next ten—much less five—years will unfold. That's not to say that it's impossible to make broad predictions about AI's impact in the coming years and decades. We've done that [elsewhere](#).

Our aim here is different: to make specific predictions about AI trends for the next 12 months, then draw out key implications for business, government, and society as a whole. We're confident in making near-term forecasts because these nascent trends are already underway, though they aren't yet attracting the attention they deserve.

We've made eight such predictions. They're based not just on insights from AI visionaries and computer scientists. They're also informed by what our leaders in assurance, consulting, and tax see on the ground with clients around the world who are grappling with how to put AI to work in their organizations and prepare their employees for a world in which AI is everywhere.

We hope you'll consider how these predictions relate to your own organization.

PwC AI predictions for 2018

1. AI will impact employers before it impacts employment
2. AI will come down to earth—and get to work
3. AI will help answer the big question about data
4. Functional specialists, not techies, will decide the AI talent race
5. Cyberattacks will be more powerful because of AI—but so will cyberdefense
6. Opening AI's black box will become a priority
7. Nations will spar over AI
8. Pressure for responsible AI won't be on tech companies alone

1

AI will impact employers before it impacts employment

Everyone has seen the headlines: *Robots and AI will destroy jobs*. But we don't see it that way. We see a more complex picture coming into focus, with AI encouraging a gradual evolution in the job market that—with the right preparation—will be positive. New jobs will offset those lost. People will still work, but they'll work more efficiently with the help of AI.

Most people have heard that AI beat the world's greatest grandmaster in chess. But not everyone knows what can usually beat an AI chess master: a "[centaur](#)," or human and AI playing chess as a team. The human receives advice from an AI partner but is also free to override it, and it's the established process between the two that is the real key to success.

This unparalleled combination will become the new normal in the workforce of the future. Consider how AI is enhancing the product design process: A human engineer defines a part's materials, desired features, and various constraints, and inputs it into an AI system, which generates a number of simulations. Engineers then either choose one of the options, or refine their inputs and ask the AI to try again.

This paradigm is one reason why AI will [strengthen the economy](#). At the same time, however, there's no denying that in some industries, economies, and roles—especially those that involve repetitive tasks—jobs will change or be eliminated. Yet in the next two years, the impact will be relatively modest: PwC's forthcoming international jobs automation study, due in February 2018, estimates that across 29 countries analyzed, the share of jobs at potential high risk of automation is only 3 percent by 2020.

Why organizations will succeed or fail

The upshot? In 2018, organizations will start realizing they need to change how they work. As they do so, they'll need to be especially mindful of what has come before: [failed tech transformations](#). There are several reasons why this happens, but two in particular are relevant to the way so many organizations are approaching AI. They're pigeon-holing AI talent. And they're thinking and working in silos.

AI-savvy employees won't just need to know how to choose the right algorithm and feed data into an AI model. They'll also have to know how to interpret the results. They'll need to know when to let the algorithm decide, and when to step in themselves.

At the same time, effective use of AI will demand collaboration among different teams. Consider an AI system that helps hospital staff decide which medical procedures to authorize. It will need input not just from medical and AI specialists, but also from legal, HR, financial, cybersecurity, and compliance teams.

Most organizations like to set boundaries by putting specific teams in charge of certain domains or projects and assigning a budget accordingly. But AI requires multidisciplinary teams to come together to solve a problem. Afterward, team members then move on to other challenges but continue to monitor and perfect the first.

With AI, as with many other digital technologies, organizations and educational institutions will have to think less about job titles, and more about tasks, skills, and mindset. That means embracing new ways of working.

A person is holding a tablet, and a stylized robot head is overlaid on the screen. The robot head has a friendly expression with large eyes and a smile.

67% of executives say AI will help humans and machines work together to be stronger using both artificial and human intelligence

Source: PwC Consumer Intelligence Series: Bot.Me, 2017
Base: 500 business executives

Implications


Popular acceptance of AI may occur quickly

As signs grow this year that the great AI jobs disruption will be a false alarm, people are likely to more readily accept AI in the workplace and society. We may hear less about robots taking our jobs, and more about robots making our jobs (and lives) easier. That in turn may lead to a faster uptake of AI than some organizations are expecting.

The organizational retooling will commence

It will be a lengthy process, but some forward-thinking organizations are already breaking down the silos that separate data into cartels and employees into isolated units. Some will also start on the massive workforce upskilling that AI and other digital technologies require. This upskilling won't just teach new skills. It will teach a new mindset that emphasizes collaboration with co-workers—and with AI.

How workers think about human-machine AI centaurs



Source: [PwC Consumer Intelligence Series: Bot.Me](#), 2017
Base: 500 business executives; percent agreeing with statement

② AI will come down to earth—and get to work

There are plenty of publications promising an AI-powered future that will look like magic: fleets of autonomous cars that never crash or encounter traffic jams, robot doctors that diagnose illness in milliseconds, and smart infrastructure that optimizes flows of people and goods and maintains itself before repairs are ever needed. All that may come—but not in 2018.

Executives think that AI will be crucial for their success: 72% believe it will be the business advantage of the future. The question is: *What can it do for me today?* And the answer is here.

Augmenting human productivity

If AI sounds far-fetched, what about a tool to perform repetitive white-collar tasks, so managers can spend their time on analysis? How about one that detects fraud and increases supply chain resilience?

This is the value of AI in 2018: it lies not in creating entire new industries (that's for the next decade), but rather in empowering current employees to add more value to existing enterprises. That empowerment is coming in three main ways:

- Automating processes too complex for older technologies
- Identifying trends in historical data to create business value
- Providing forward-looking intelligence to strengthen human decisions

Value from tedious tasks

Consider how most companies' finance functions spend a large portion of their time: wading through data coming from ERP, payment processing, business intelligence, and other systems. Many staff members also spend hours each day poring through legal contracts and emails, or performing mundane transactional tasks.

The result is that value-adding analysis is what many finance professionals only do when they have time left over from their other, routine tasks.

Now imagine an AI system scanning all the function's data, identifying trends and anomalies, performing many transactions automatically, and flagging relevant issues for further attention. Imagine AI also identifying and explaining likely risks and offering data-driven forecasts to support managers' analysis and decisions.

It may not be as sexy as a smart city, but this kind of **practical AI** is ready right now. And it's often "sneaking in through the backdoor." Enterprise application suites from Salesforce, SAP, Workday, and others are increasingly incorporating AI.

Where industries will put practical AI to work

Ranking of AI impact by its potential to free up time, enhance quality, and enhance personalization

Ranking	Industry	High-potential use cases
 1	Healthcare	- Supporting diagnosis by detecting variations in patient data - Early identification of potential pandemics - Imaging diagnostics
 1	Automotive	- Autonomous fleets for ride sharing - Semi-autonomous features such as driver assist - Engine monitoring and predictive, autonomous maintenance
 3	Financial services	- Personalized financial planning - Fraud detection and anti-money laundering - Automation of customer operations
 4	Transportation and logistics	- Autonomous trucking and delivery - Traffic control and reduced congestion - Enhanced security
 5	Technology, media, and telecommunications	- Media archiving, search, and recommendations - Customized content creation - Personalized marketing and advertising
 6	Retail and consumer	- Personalized design and production - Anticipating customer demand - Inventory and delivery management
 7	Energy	- Smart metering - More efficient grid operation and storage - Predictive infrastructure maintenance
 8	Manufacturing	- Enhanced monitoring and auto-correction of processes - Supply chain and production optimization - On-demand production

Source: [PwC Global AI Impact Index](#), 2017

Implications

Business problems will open the door to AI

Leaders don't need to adopt AI for AI's sake. Instead, when they look for the best solution to a business need, AI will increasingly play a role. Does the organization want to automate billing, general accounting and budgeting, and many compliance functions? How about automating parts of procurement, logistics, and customer care? AI will likely be a part of the solution, whether or not users even perceive it.

New kinds of ROI measures are needed

Sometimes the best approach to gauge AI's value is to use the same measures you'd apply to any other business investment: metrics such as increased revenue or reduced costs. But AI's most powerful benefits are often indirect, so organizations will want to explore other measures of ROI. Automated full-time equivalents can capture how AI is freeing human workers from mundane tasks. Other metrics can show how AI is improving human decision-making and forecasts.



54% of business executives say AI solutions implemented in their businesses have already increased productivity


Source: [PwC Consumer Intelligence Series: Bot.Me, 2017](#)
Base: 500 business executives

3

AI will help answer the big question about data

Many companies haven't seen the payoff from their big data investments. There was a disconnect. Business and tech executives thought they could do a lot more with their data, but the learning curve was steep, tools were immature, and they faced considerable organizational challenges.

Few businesses get value from their data—but AI could change that



Source: [PwC 2017 Global Digital IQ Survey](#)

Q: To what extent do you agree with the following statement (strongly agree): We effectively utilize all the data we capture to drive business value

Bases: 72, 217, 135, 322, 237, 75, 375, 131, 156, 433

Now, some are rethinking their data strategy as the landscape matures and AI itself becomes more real and practical. They're starting to ask the right questions, like: *How can we make our processes more efficient?* and *What do we need to do to automate data extraction?*

At the same time, organizations are now able to take advantage of new tools and technical advancements, including:

- Easier methods for mining less-structured data, including natural language processing for text indexing and classification
- Enterprise application suites that incorporate more AI
- Emerging data lake-as-a-service platforms
- Public clouds that can take advantage of different kinds of data
- Automated machine learning and data management

Feeding the AI beast

Despite these advances, many organizations still face a challenge. Many kinds of AI, such as supervised machine learning and deep learning, need an enormous amount of data that is standardized, labeled, and “cleansed” of bias and anomalies. Otherwise, following the ancient rule—*garbage in, garbage out*—incomplete or biased data sets will lead to flawed results. The data must also be specific enough to be useful, yet protect individuals’ privacy.

Consider a typical bank. Its various divisions (such as retail, credit card, and brokerage) have their own sets of client data. In each division, the different departments (such as marketing, account creation, and customer service) also have their own data in their own formats. An AI system could, for example, identify the bank’s most profitable clients and offer suggestions on how to find and win more clients like them. But to do that, the system needs access to the various divisions’ and departments’ data in standardized, bias-free form.

The right approach to data

It’s rarely a good idea to *start* with a decision to clean up data. It’s almost always better to start with a business case and then evaluate options for how to achieve success in that specific case.

A healthcare provider, for example, might aim to improve patient outcomes. Before beginning to develop the system, the provider would quantify the benefits that AI can bring. The provider would next look at what data was needed—electronic medical records, relevant journal

articles, and clinical trials data, among others—and the costs of acquiring and cleansing this data.

Only if the benefits—including measures of indirect benefits and how future applications can use this data—exceed the costs should this provider move forward.

That's how many organizations will ultimately reform data architecture and governance: with AI and other technologies offering value propositions that require it.

Implications

Success will lead to success

Those enterprises that have already addressed data governance for one application will have a head start on the next initiative. They'll be on their way to developing best practices for effectively leveraging their data resources and working across organizational boundaries.

Third-party data providers will prosper

There's no substitute for organizations getting their internal data ready to support AI and other innovations, but there is a supplement: Vendors are increasingly taking public sources of data, organizing it into data lakes, and preparing it for AI to use.

More synthetics are coming

As data becomes more valuable, advances in synthetic data and other “lean” and “augmented” data learning techniques will accelerate. We may not need, for example, a whole fleet of autonomous cars on the road to generate data about how they'll act. A few cars, plus sophisticated mathematics, will be sufficient.



59% of executives say big data at their company would be improved through the use of AI

Source: PwC Consumer Intelligence Series: Bot.Me, 2017
Base: 500 business executives

4

Functional specialists, not techies, will decide the AI talent race

As AI spreads into more specific areas, it will require knowledge and skill sets that data scientists and AI specialists usually lack.

Consider a team of computer scientists creating an AI application to support asset management decisions. The AI specialists probably aren't experts on the markets. They'll need economists, analysts, and traders working at their side to identify where the AI can best support the human asset manager, help design and train the AI to provide that support, and be willing and able to use the AI effectively.

And since the financial world is in constant flux, once the AI is up and running, it will need continual customizing and tweaking. For that too, functional specialists—not programmers—will have to lead the way. The same is true not just throughout financial services, but in healthcare, retail, manufacturing, and every sector that AI touches.

Citizen data scientists wanted

AI is becoming more user friendly. Users no longer need to know how to write code in order to work with some AI applications. But most still demand far more technical knowledge than a spreadsheet or word processing program does.

For example, many AI tools require users to formulate their needs into machine learning problem sets. They also require an understanding of which algorithms will work best for a particular problem and a particular data set.

The exact level of knowledge required will vary, but we can broadly divide AI's demands on human knowledge into three categories. First, most members of an AI-enabled enterprise will need some basic knowledge of AI's value as well as what it can and can't do with data. Second, even the most mature AI program will always need a small team of computer scientists. The third group, which many organizations aren't yet paying attention to, are AI-savvy functional specialists.

They won't have to be programmers. They will have to understand the basics of data science and data visualization and the basics of how AI "thinks." They'll have to be citizen data scientists.

Retail analysts, engineers, accountants, and many other domain experts who know how to prepare and contextualize data so AI can make optimal use of it will be crucial to enterprise success. As AI leaves the computer lab and enters everyday work processes, these functional specialists will be even more important than computer scientists.

Implications

Faster upskilling means faster AI deployment

Enterprises that intend to take full advantage of AI shouldn't just bid for the most brilliant computer scientists. If they want to get AI up and running quickly, they should move to provide functional specialists with AI literacy. Larger organizations should prioritize by determining where AI is likely to disrupt operations first and start upskilling there.

Upskilling will lead to new approaches to learning

Organizations will have to upskill many of their employees to learn the basics of data science and how to think like an AI application. Given the enormity of this task, companies must find ways to assess the skills of high-potential learners and put them on individual learning paths, to get them up to speed quickly.



67% of jobs requiring data science and analytics skills are in fields other than AI

Source: PwC and Business Higher Education Forum, [Investing in America's data science and analytics talent, 2017](#)

5 Cyberattacks will be more powerful because of AI—but so will cyberdefense

What's one job where AI has already shown superiority over human beings? Hacking. Machine learning, for example, can easily enable a malicious actor to follow your behavior on social media, then customize phishing tweets or emails—just for you. A human hacker can't do the job nearly as well or as quickly.

The more AI advances, the more its potential for cyberattacks grows too. Techniques like advanced machine learning, deep learning, and neural networks enable computers to find and interpret patterns. They can also find and exploit vulnerabilities.

Intelligent malware and ransomware that learns as it spreads, machine intelligence coordinating global cyberattacks, advanced data analytics to customize attacks—unfortunately, it's all on its way to your organization soon. And AI itself, if not well-protected, gives rise to new vulnerabilities. Malicious actors could, for example, inject biased data into algorithms' training sets.

AI to the rescue

Just as we expect AI to be a growing cyberthreat this year, we're also confident it will be part of the solution. Already, [scalable machine learning techniques](#) combined with cloud technology are analyzing enormous amount of data and powering real-time threat detection and analysis. AI capabilities can also quickly identify “hot spots” where cyberattacks are surging and provide cybersecurity intelligence reports.

The winner of the US Defense Department's [DARPA Cyber Grand Challenge](#), a cybersecurity competition, used AI deep learning—and the [Pentagon has purchased](#) the technology.

Yet even in cybersecurity, some things only people can do. Humans are better at absorbing context and thinking imaginatively. Cyberwars won't simply be two sets of computers battling it out. But AI will become an important part of every major organization's cybersecurity toolkit.

Implications

Don't bring a knife to a gunfight

In other parts of the enterprise, many organizations may choose to go slow on AI, but in cybersecurity there's no holding back: Attackers will use AI, so defenders will have to use it too. If an organization's IT department or cybersecurity provider isn't already using AI, it has to start thinking immediately about AI's short- and long-term security applications. Sample use cases include distributed denial of service (DDoS) pattern recognition, prioritization of log alerts for escalation and investigation, and risk-based authentication.

Cybersecurity may speed up AI's acceptance

Since even AI-wary organizations will have to use AI for cybersecurity, cyberdefense will be many enterprises' first experience with AI. We see this fostering familiarity with AI and willingness to use it elsewhere. A further spur to AI acceptance will come from its hunger for data: The greater AI's presence and access to data throughout an organization, the better it can defend against cyberthreats. Some organizations are already building out on-premise and cloud-based "threat lakes," that will enable AI capabilities.

An AI hack may increase public fears

Many people are already nervous about AI. Even more are concerned about cybersecurity. It's possible that when AI makes headlines this year, it won't be for helping humanity—but for having enabled a major hack. Better cybersecurity can mitigate this risk. Besides developing feature sets for AI capabilities, this heightened security will require companies to augment the data and compute platforms that support advanced analytics with privileged access monitoring, object-level change management, source code reviews, and expanded cybersecurity controls, among other precautions.



27% of executives say their organization plans to invest this year in cybersecurity safeguards that use AI and machine learning

Source: PwC 2018 Global State of Information Security® Survey
Base: 9,500 business and technology executives

6 Opening AI's black box will become a priority

Might AI-powered autonomous weapons become serial killers? Could an AI system told to reduce air pollution decide that the most logical way to do so is to eliminate the human race? Such fears may make for good thrillers, but the danger is manageable.

Here's the secret about AI that many of its proponents don't like to mention: It's not that smart—at least not yet. AI is getting better at pattern and image recognition, automating complex tasks, and helping humans make decisions. All that offers opportunities for enterprises that could be worth [trillions of dollars](#).

In the past, for example, to teach an AI program chess or another game, scientists had to feed it data from as many past games as they could find. Now they simply provide the AI with the game's rules. In a few hours it figures out on its own how to beat the world's greatest grandmasters.

That's extraordinary progress, with immense potential to support human decision making. Instead of playing chess, an AI program with the right rules can “play” at corporate strategy, consumer retention, or designing a new product.

But it's still just following rules that humans have devised. With appropriate attention paid to [responsible AI](#), we can safely harness its power.

A real risk


If AI should always be controllable, it's not always understandable. Many AI algorithms are beyond human comprehension. And some AI vendors will not reveal how their programs work to protect intellectual property. In both cases, when AI produces a decision, its end users won't know how it arrived there. Its functioning is a “black box.” We can't see inside it.

That's not always a problem. If an ecommerce website uses mysterious algorithms to suggest a new shirt to a consumer, the risks involved are low.

But what happens when AI-powered software turns down a mortgage application for reasons that the bank can't explain? What if AI flags a certain category of individual at airport security with no apparent justification? How about when an AI trader, for mysterious reasons, makes a leveraged bet on the stock market?

Users may not trust AI if they can't understand how it works. Leaders may not invest in AI if they can't see evidence of how it made its decisions. So AI running on black boxes may meet a wave of distrust that limits its use.

What it means to look inside AI's black box



Source: PwC

Implications

Many black boxes will open

We expect organizations to face growing pressure from end users and regulators to deploy AI that is explainable, transparent, and provable. That may require vendors to share some secrets. It may also require users of deep learning and other advanced AI to deploy new techniques that can explain previously incomprehensible AI.

Organizations face tradeoffs

Most AI can be made explainable—but at a cost. As with any other process, if every step must be documented and explained, the process becomes slower and may be more expensive. But opening black boxes will reduce certain risks and help establish stakeholder trust.

Enterprises need a framework for AI explainability decisions

Explainability, transparency, and provability aren't absolutes; they exist on a scale. A framework to assess business, performance, regulatory, and reputational concerns can enable optimal decisions about where each AI use case should fall on that scale. A healthcare firm using AI to help make life-or-death decisions has different needs than a private equity fund using AI to identify potential targets for further research.

7 Nations will spar over AI

AI is going to be big: \$15.7 trillion big by 2030, according to [our research](#). The AI pie is so big that besides individual companies, countries are working on strategies to claim the biggest possible slice.

The US started off strong, with a [trio of reports](#) in 2016. They outlined a plan to make the US an AI powerhouse and thereby boost both the economy and national security.

Recommendations included increased federal funding, regulatory changes, the creation of shared public data sets and environments, the definition of standards and benchmarks, workforce development, and ways for AI to bolster cybersecurity and the military.

But since a new administration entered at the start of 2017, the government has abandoned this plan. It's cutting AI research funds.

Its recently passed tax reform, however, could boost AI in the US. The lower corporate tax rate, provisions for repatriating cash from overseas, and permission to expense 100 percent of capital investments is likely to spur investment in AI and other technologies. And the current administration's emphasis on deregulation could help AI in certain sectors, such as drones and autonomous vehicles.

The new AI leaders

While the US government moves on from its AI development plan, other countries are taking action.

The [UK](#) last year launched a plan to improve access to data, AI skills, and AI research and uptake. Its latest [budget](#) added funding for a Centre for Data Ethics and Innovation to drive responsible AI, for exploratory work on data trusts, and for new AI fellowships and researchers.

Canada—already an AI leader—is also working to make AI a lynchpin of its future economy. The federal government last year launched its [Pan-Canadian Artificial Intelligence Strategy](#). The program includes funding for AI research centers in collaboration with private companies and universities. It also aims to attract and retain top AI talent.

Japan meanwhile released an [AI technology strategy](#) with a three-phase plan to achieve a true AI ecosystem. Building on successes in robotics, Japan's government envisions joining AI with other advanced technologies, such as the internet of things, autonomous vehicles, and the blending of cyber and physical space.

Other countries with recently released national AI plans include [Germany](#) with ethical guidelines for automated driving and its [Industrie 4.0](#) initiative, and the [UAE](#) with a strategy to use AI to boost government performance and various economic sectors.

China stands apart

China's next-generation [AI plan](#), released in 2017, declared AI as a strategic national priority for the country and showcased the top leadership's vision for a new economic model driven by AI.


Unlike the US, the Chinese government is putting this plan into practice. For example, it recently commissioned Baidu to create a national "deep learning laboratory" together with leading universities—and it's investing an undisclosed sum in the effort.

The country is already strong in AI. Baidu, Alibaba, and Tencent are among the global AI leaders. Chinese programmers recently won the ImageNet AI competition. And its leading ecommerce companies are using highly sophisticated AI in their warehouses and across the business.

Other countries also have innovative engineers, universities, and companies. But China stands apart in the extent to which its government is prioritizing AI. Our [research](#) indicates that China will reap the most benefit from AI over the next decade: some \$7 trillion in GDP gains by 2030, thanks to an uptick in productivity and consumption.

Where AI gains will be realized

AI's impact on GDP by 2030



Source: [PwC Global Artificial Intelligence Study](#), 2017

Implications

China's investment may awaken the West

If China starts to produce leading AI developments, the West may respond. Whether it's a "Sputnik moment" or a more gradual realization that they're losing their lead, policymakers may feel pressure to change regulations and provide funding for AI.

More national and regional strategies will come

More countries should issue AI strategies, with implications for companies. It wouldn't surprise us to see Europe, which is already moving to protect individuals' data through its [General Data Protection Regulation \(GDPR\)](#), issue policies to foster AI in the region.

Collaboration may come too


National competition for AI will never cease—there's too much money at stake. But we do expect growing opportunities, facilitated by the UN, the World Economic Forum, and other multilateral organizations, for countries to collaborate on AI research in areas of international concern.

8 Pressure for responsible AI won't be on tech companies alone

New technologies often bring new fears, justified or not, and not just among conspiracy theorists. Seventy-seven percent of CEOs in a 2017 [PwC survey](#) said AI and automation will increase vulnerability and disruption to the way they do business. Odds are good that if we asked government officials, the response would be similar.

Leaders will soon have to answer tough questions about AI. It may be community groups and voters worried about bias. It may be clients fearful about reliability. Or it may be boards of directors concerned about risk management, ROI, and the brand.

What's holding AI back in the enterprise?



Source: [PwC CEO Pulse Survey](#), 2017

Q: Which of the following issues surrounding AI adoption concern you the most?

Base: 239

In all cases, stakeholders will want to know that organizations are using AI [responsibly](#), so that it strengthens the business and society as a whole.

The result, we believe, will be pressure to adopt principles for responsible AI.

A global movement begins

We're not alone in this belief. The World Economic Forum's [Center for the Fourth Industrial Revolution](#), the [IEEE](#), [AI Now](#), [The Partnership on AI](#), [Future of Life](#), [AI for Good](#), and [DeepMind](#), among other groups, have all released sets of principles that look at the big picture: how to maximize AI's benefits for humanity and limit its risks.

Some areas of relative consensus (which we fully support) among these institutions include:

- **designing** AI with an eye to societal impact
- **testing** AI extensively before release
- **using** AI transparently
- **monitoring** AI rigorously after release
- **fostering** workforce training and retraining
- **protecting** data privacy
- **defining** standards for the provenance, use, and securing of data sets
- **establishing** tools and standards for auditing algorithms

With any new technology (and many old ones too), the golden rule we follow is to do more than compliance requirements demand. Regulators and laws often lag innovation. Organizations that don't wait for policymakers to issue orders, but instead voluntarily use new technology responsibly, will reduce risks, improve ROI, and strengthen their brands.

Implications

New structures for responsible AI

As organizations face pressure to design, build, and deploy AI systems that deserve trust and inspire it, many will establish teams and processes to look for bias in data and models and closely monitor ways malicious actors could “trick” algorithms. Governance boards for AI may also be appropriate for many enterprises.

Public-private partnerships and public-citizen partnerships

One of the best ways to use AI responsibly is for public and private sector institutions to collaborate, especially when it comes to AI's societal impact. Likewise, as more governments explore the use of AI to distribute services efficiently, they're engaging citizens in the process. In the UK, for example, the RSA (Royal Society for the encouragement of Arts, Manufactures and Commerce) is conducting a series of [citizen juries](#) on the use of AI and ethics in criminal justice and democratic debate.

Self-regulatory organizations to facilitate responsible innovation

Since regulators may scramble to keep up, and self-regulation has its limits, [self-regulatory organizations](#) (SROs) may take the lead with responsible AI. An SRO would bring users of AI together around certain principles, then oversee and regulate compliance, levy fines as needed, and refer violations to regulators. It's a model that has worked in other industries. It may well do the same for AI and other technologies.

Authors

Mike Baccala

US Assurance Innovation
Leader

Chris Curran

US New Ventures Chief
Technology Officer

Dan Garrett

US Technology Consulting
Leader

Scott Likens

New Services and
Emerging Technologies
Leader, US, China and
Japan

Anand Rao

Global and US Artificial
Intelligence Leader

Andy Ruggles

US Tax Reporting and
Strategy Leader

Michael Shehab

US Tax Technology and
Process Leader

Editor

Chrisie Wendin

US Editorial Director,
Technology

Marketer

Sarah Weiss

US Tech and Emerging
Technology Marketing

Contributors

Cristina Ampil

US Integrated Content,
Managing Director

Maria Axente

UK AI Programme Driver

Bjarne Berg

US Tax Analytics and
Innovation Leader

Richard Berriman

UK Strategy & Data Science
Innovation Team

Euan Cameron

UK AI Leader

Chris Castelli

US Editorial Director,
Risk and Regulatory

Mike Flynn

US Assurance Principal

Ilana Golbin

US AI Accelerator Manager

John Hawksworth

UK Chief Economist

Mir Kashifuddin

US Cybersecurity and Privacy
Principal

Art Kleiner

Editor in Chief, PwC Global

Matt Lieberman

US Advisory Marketing Leader

Rob McCargow

UK AI Programme Leader

Alan Morrison

US Integrated Content,
Technology

Diego Muro

US Tax Transfer Pricing
Principal

Horacio Pena

US Tax Transfer Pricing Leader

Tom Puthiyamadam

Global Digital Services Leader

Pia Ramchandani

US AI Accelerator Director

David Resseguie

US Emerging Tech Tax Leader,
Tax Technology

David Sapin

US Advisory Risk and
Regulatory Leader

Jas Sidhu

UK Disruptive Innovation Lead

John Sviokla

US Marketing Leader

Gerard Verweij

Global Data & Analytics
Leader