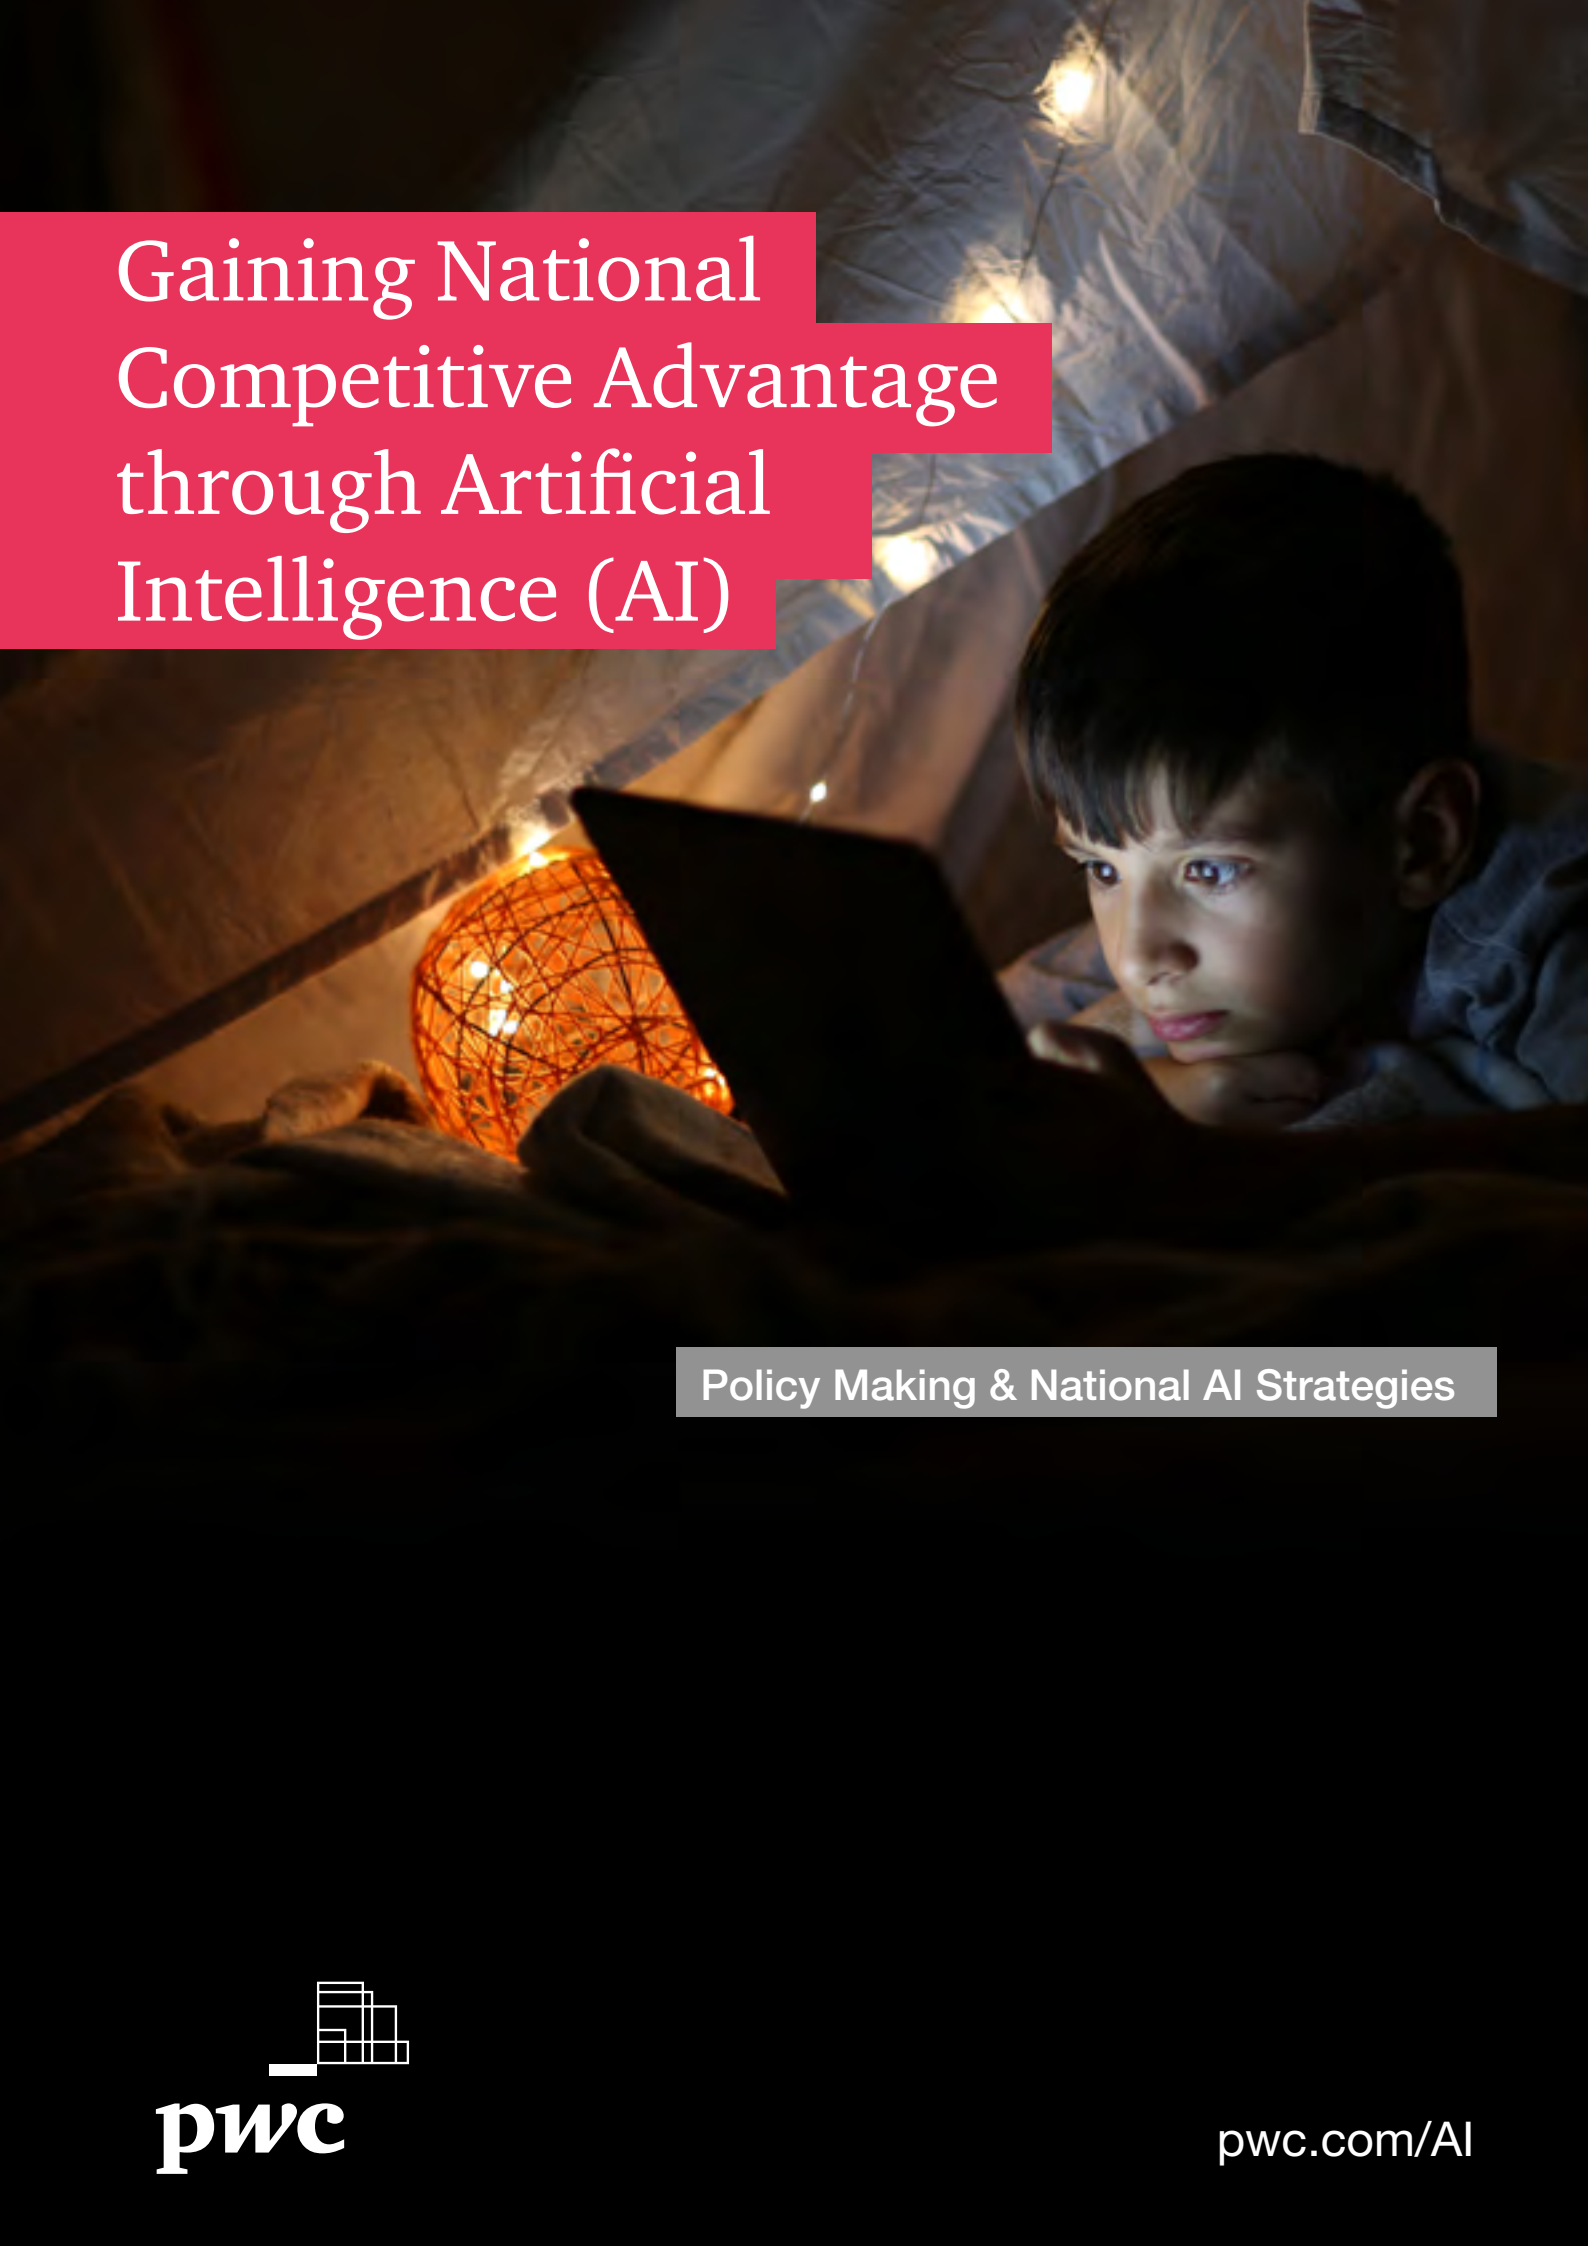

Gaining National Competitive Advantage through Artificial Intelligence (AI)

Policy Making & National AI Strategies

pwc

pwc.com/AI

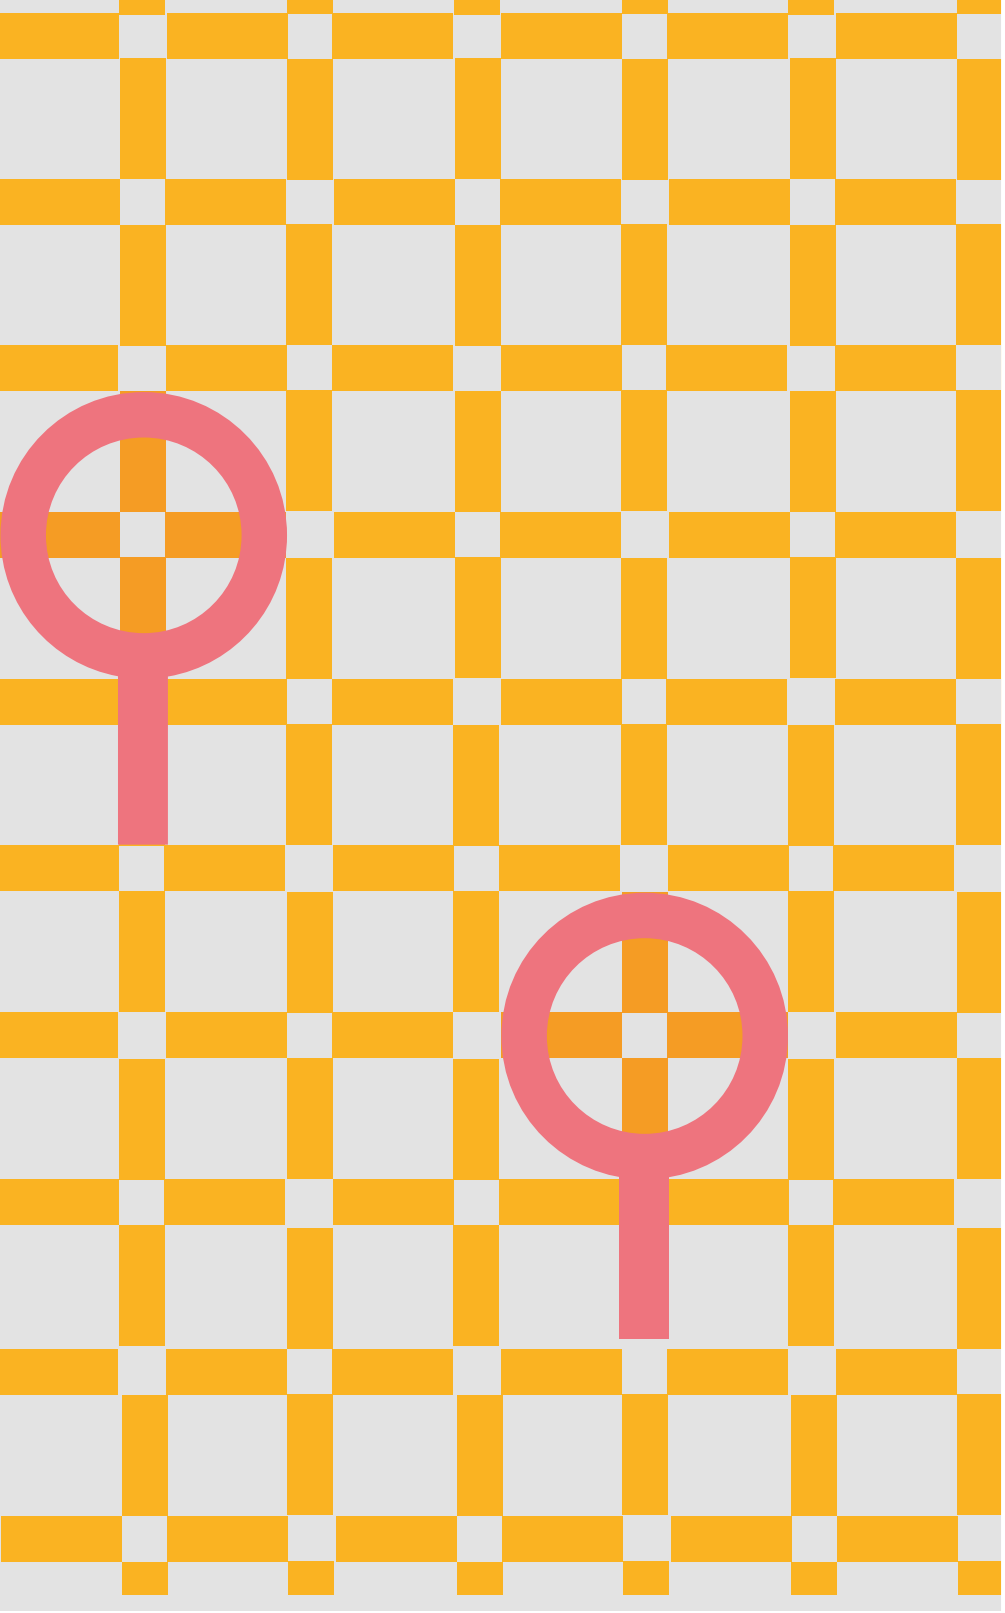

Table of contents

Introduction

Governments are vying with each other to develop national AI strategies to attract and foster business investment and innovation; educate, train, and create a skilled workforce now and into the future; provide the benefits of AI to their citizens, while protecting them from some of the potential risks of AI. Cognitive, behavioural, and network capital are the three primary sources of capital that governments need to leverage to gain competitive advantage. AI policy making requires a number of trade-offs that will ultimately be driven by societal values and what each nation wants. Governments can make these trade-offs, with businesses, consumer advocacy groups, and international bodies actively involved in shaping public opinion. These choices can determine whether AI contributes to the betterment of humanity.

Artificial intelligence (AI) presents vast opportunities, but is not without potential pitfalls and risks. This paradox has become increasingly evident for government leaders. They want to give domestic companies an edge over the competition, but are also expected to protect their citizens and use AI for social good. They want to support innovation, while still maintaining some level of control over how new technologies impact society at large. With a huge payoff potentially on the line, AI has the potential to increase worldwide GDP by 14 percent by 2030, an infusion of US$15.7 trillion into the global economy [Rao et al, 2017] — it should come as no surprise that governments are eager to claim their share.

To date, more than 20 countries and regions, including Canada, China, France, Germany, India, Japan, Kenya, Mexico, New Zealand, Russia, South Korea, the United Arab Emirates, and the U.K., have released AI strategy documents. Global bodies such as the World Economic Forum and industry associations such as the Partnership on AI are convening committees and acting in an advisory capacity in some cases. The Institute of Electrical and Electronics Engineers has also released standards for ethical AI design.

Overall, these new policies outline how governments plan to foster AI development to encourage domestic companies to develop solutions that will boost gross domestic product (GDP) and offer a host of societal benefits. At the same time, they tackle questions about security, privacy, transparency, and ethics. Given the potential for AI to have disruptive social and environmental effects, the development of sophisticated national and international governance structures will become increasingly critical. Perhaps no other emerging technology has inspired such scrutiny and discussion.

Such activity introduces an imperative for business leaders to look for ways to help shape and refine the national AI strategies that will impact the regions in which they operate. PwC research reveals that companies recognise this need. PwC's 22nd Global CEO Survey [PwC, 2019] found that 85 percent of CEOs agree that AI will significantly change the way they do business in the next five years. The survey also found that CEOs believe that AI is good for society — and more than two-thirds of CEOs agree that governments should play a critical and integral role in AI development.

In a recent survey of over 1,000 companies, conducted for PwC U.S. AI Predictions 2019, 37 percent of executives rated "ensuring AI systems are trustworthy" as their top priority (See Figure 1).

Figure 1: Top five AI challenges for 2019

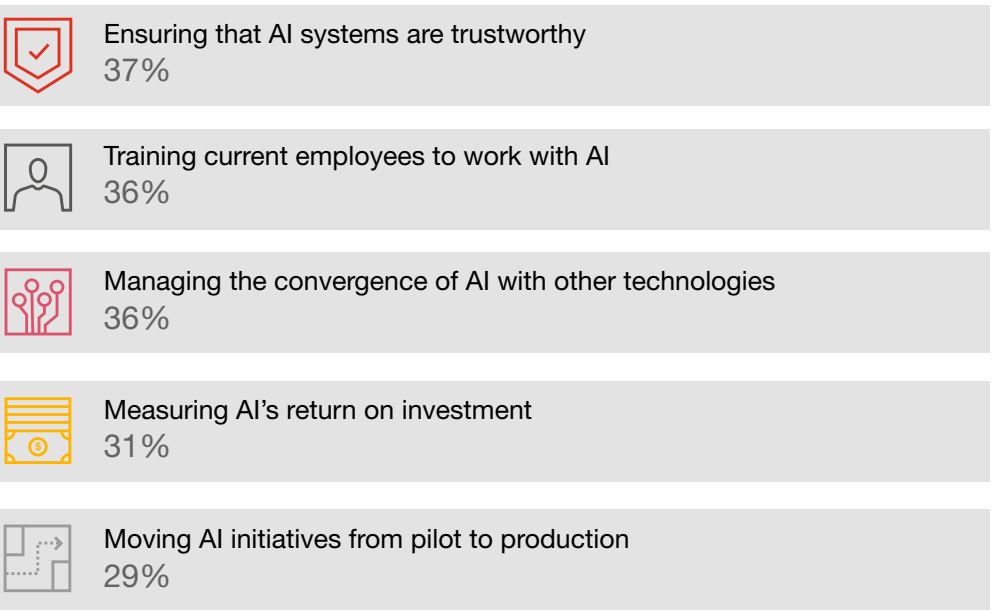

Source: PwC 2019 AI Predictions
Base: 1,001
Q: Which AI data-related issues will be the top priorities for your organisation in 2019?

1 From a Bionic Company to a Bionic Nation

Given the exponential growth of technology it is understandable that nations are vying with each other to develop their own national AI strategy. But how can AI advance a nation's competitive advantage? Delving into the six different forms of capital bionic companies use to build competitive advantage [Everson and Sviokla, 2018], and applying them to nations, will help answer this question.

Right up to the industrial revolution of the nineteenth century, the main competitive advantage of nations was **natural capital** (value of land, water, and other environmental resources) and **human capital** (the physical and intellectual ability of the rulers, traders, and people within the country). However, as the world moved into the industrial revolution during the nineteenth and twentieth centuries **financial capital** also started playing a significant role with a rise in different mechanisms for raising financial capital and distributing it to where it was needed. This led to the growth of the stock exchanges and the Central Bank institutions to manage the supply and demand of money. National governments controlled these institutions until the recent past. Stock exchanges started getting privatised toward the late twentieth and early twenty-first centuries. Central Banks are still under the direction of national governments with varying levels of independence from political influence.

With the advent of the information revolution in the late nineteenth century, the world started seeing a new form of capital — **behavioural capital** — emerge. Behavioural capital is the collection and modeling of data that tracks the behaviour of people, companies, nature, and manufactured things [Everson 2018]. While the information revolution started the creation of behavioural capital in service and manufacturing sectors, for example, the creation of customer records by banks, the bill of materials for manufacturing, the advent of the internet revolution has accelerated the creation of this behavioural capital. From large organisations the creation of this behavioural capital has moved to small businesses and consumers. The growth of social media has resulted in capturing the behaviour of consumers. With the internet of things (IoT) and industrial IoT there may be billions of physical assets with built-in sensors, leading to a more comprehensive behaviour capital of "things", in addition to people and companies.

Traditionally, this behavioural capital has resided with companies offering certain services to their customers, for example, internet search behaviour, financial purchase behaviour, social media behaviour, and online shopping behaviour. Nations are now coming to terms with the value of this behavioural capital and the potential risks of how the behaviours of their citizens can be first understood, then manipulated, leading to an increased interest in gaining control over and regulating the use of behavioural capital.

The advent of the internet revolution in the early twenty-first century quickly accelerated the creation of **network capital**. Network capital is the creation and management of nodes, and interactions between these nodes. The telephone network of the twentieth century is arguably the first global network capital. However, the network effect of the internet, when the number of websites grew from 3.1 million in year 1999 to 38.7 million in year 2002 [JWYG, 2017] really tells the story of network capital. Once again, national governments have been somewhat slow to grasp the potential impact of such technologies on their citizens as well as the businesses based in their countries.

With the emergence of the AI revolution — the collection of analytics, automation, and artificial intelligence technologies — emerges **cognitive capital**. Cognitive capital is the set of algorithms, knowledge, and insights synthesised from the information and decision flows between individuals and companies. For example, AlphaGo, a computer program that beat a human world Go champion, has all the knowledge, experience, and insights required to beat any human Go player. In other words, it has the cognitive capital for Go. It is not just games where companies are building cognitive capital, companies

are building **cognitive capital** for hedge funds, for claims, for medical diagnoses and much more. The ability to capture cognitive capital is now moving from individual companies to national governments. Just imagine the social good a government could provide its citizens if it could build the cognitive capital for the health of all citizens — not just how diseases are diagnosed and treated, but all also devising behavioural interventions to improve citizens' overall well-being.

All six forms of capital are interrelated, with each reinforcing the others. For example, rideshare and ride-hail companies are starting to collect detailed traffic data on the demand and supply patterns of rideshare and ride-hail in many major cities. This rideshare/ride-trail driving behaviour is a microcosm of the overall traffic patterns in these cities. Rideshare/ride-hail is also a sector that benefits from network effects. The more vehicles that are in the rideshare/ride-hail fleet, the more beneficial it is for individuals to use or be a member of the fleet. By building AI algorithms that can synthesise and project these driving behaviours over time, the company (or others in the sector), could start developing the cognitive capital for transportation, congestion, and pollution for these major cities. They could also start monetising this cognitive capital, creating more financial capital for themselves. Governments could then use this cognitive capital for city planning, sustainable development, and so on.

While the six forms of capital can be exploited to build a competitive advantage in AI at the national level, if the activity is not managed properly, nations could also see their competitive advantage decline due to some of the risks of AI. The next section looks at the possible negative side effects of AI and outlines some of the key risks.

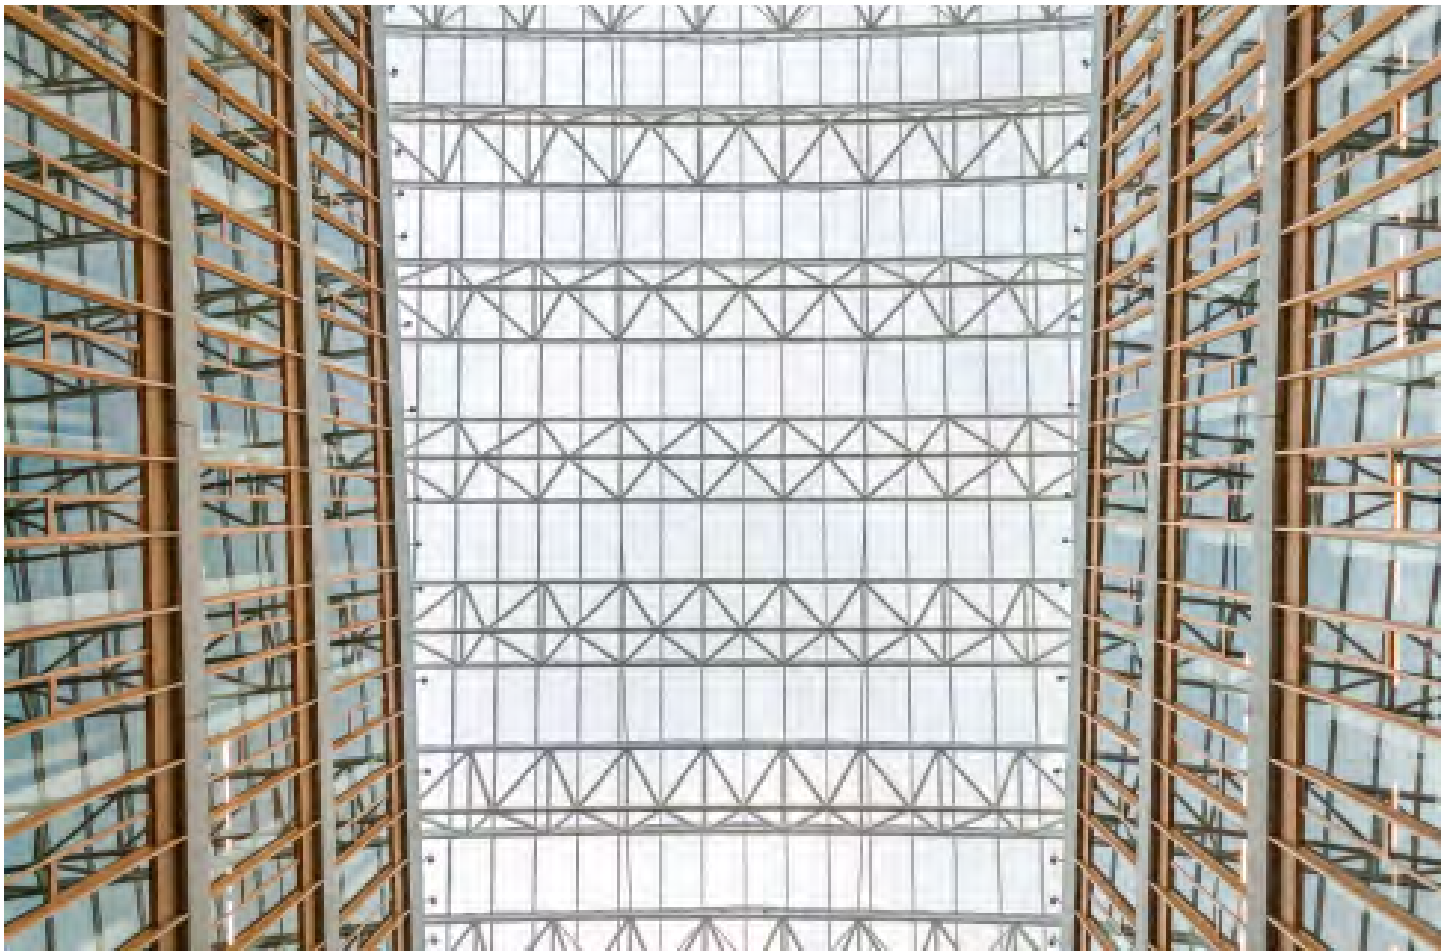

2 Risks of Artificial Intelligence

National AI policies have significant ground to cover. Besides working to increase domestic competitiveness and help businesses succeed with AI using the six forms of capital identified above, these policies also aim to address certain key concerns that accompany the technology. Companies that develop AI applications often face the same concerns, but governments can offer a model for businesses to follow while helping address some of the risks of AI.

At PwC, we group the risks associated with AI into six types. The first three categories of risks — social, ethical, and economic — are more important, far-reaching, and impact the socio-economic environment in which AI systems operate. These are issues that need to be tackled at the national, supra-national, or societal level, across the globe. The next three categories of risks — performance, control, and security — are related to how AI systems are scoped, designed, built, tested, monitored, and refined. These risks occur at the enterprise or business-unit level. Mitigating these risks may require changes to the way organisations build and govern AI systems. The issues related to social, ethical and economic risks are of a higher order significance than the technical issues related to the other three risks - these latter must be managed and resolved in order to support society and humanity, and should be subordinate to this purpose.

Social risks

Large-scale automation threatens to reduce employment in the transportation, manufacturing, agriculture and the service sectors, among others. Higher unemployment rates could lead to greater inequality in society. In addition, algorithms designed by a subset of the population at a national and global level have the potential for unconscious bias, possibly further marginalising minorities or other groups. Autonomous weapons also pose a significant threat to society, possibly creating bigger, more dangerous conflicts faster. Once unleashed, this might lead to rapid and significant environmental damage, even a "doomsday" scenario where weaponised AI threatens humanity's existence.

Ethical risks

The ethical and responsible use of AI involves three main elements: the use of big data; the growing reliance on algorithms to perform tasks, shape choices and make decisions; and the gradual reduction of human involvement in many processes. Together, these elements raise issues related to fairness, accountability, equality and respect for human rights. Additionally, while biased AI outcomes can raise significant privacy concerns, many insights and decisions about individuals are based on inferred group or community attributes. Accordingly, consideration of the harm AI could do must be framed beyond the individual level. It must be also recognised that privacy is not the only issue.

Figure 2 provides a sample of specific risks for each of these categories. While not claiming to be exhaustive, the framework provides a classification that can be further enhanced as more risks are identified.

Economic risks

As companies adopt AI, it may alter the competitive landscape, creating winners and losers. Those able to improve their decision-making most quickly through AI may find the benefits accelerate very quickly, while slower adopters may be left behind. Companies that struggle during the AI transition may be forced to reduce their investment in AI, possibly impairing their profitability and potentially threatening their own existence. Given the potential for accelerating returns on the cognitive capital (a combination of human and machine intelligence), the first movers with the right data and experts can quickly monopolise their market. Given the global nature of the digital world, this could very quickly result in a global race for supremacy, forcing governments to intervene to protect their local industries and potentially paving the way for more protectionism and less globalisation.

Performance risks

Like any other software system, AI systems need to be verified and validated using standard methodologies. However, AI systems — particularly machine learning systems — differ significantly from standard software systems. Broadly there are two phases to building a machine learning system [Dietterich 1988; Hall et al, 2017]. First, the developer trains the system on a particular task (such as classifying data) by providing large volumes of input data as well as output data or labels. Once the system has been adequately trained, it is deployed in a production mode, where, given a new piece of data, the machine learning system can predict the output (for example, classification of the new data).

Unlike traditional systems, machine learning systems cannot be verified by analysing code line by line. Instead, there is a need to ensure that the data provided is representative, there is no bias in the data, and it is understood how the system is identifying the features and how it is making the recommendations. The difficulty of accomplishing this for many machine learning algorithms makes them a "black box," meaning it is not easy to ascertain whether the performance or outputs of AI algorithms are accurate or desirable. The emerging field of explainable AI (XAI) research [Gunning 2016] aims to create new AI methods that can explain their inner workings.

Another major concern with respect to AI algorithms is the potential for these algorithms to institutionalise bias. Machine learning algorithms use historical data to detect patterns and make inferences. Thus, using historical data, even if it is factual, can lead to biased outcomes. For example, to see this in action, just do an internet image search for the terms nurse and doctor and you'll see certain gender stereotypes emerge. The machine learning algorithm will tend to conclude that nurses are generally female and doctors are generally male. Such bias must be mitigated when it can lead to discrimination against a particular group of people. However, as shown elsewhere [Rao and Golbin, 2019] the notion of fairness is a social construct and there are many different mathematical definitions of fairness. It is impossible for an AI algorithm to be considered universally "fair" across all of these definitions.

But this field is still in its early days. Meanwhile, ongoing research aims to reduce model bias resulting from biases in training data, while also increasing the stability of model performance. As AI solutions are deployed, one unintended consequence could be the overreliance on AI algorithms with variable performance [Goodman and Flaxman, 2016]. It is essential that humans stay "in the loop" to audit algorithm outputs to help mitigate these unintended biases and wider performance risks.

Control risks

Some AI systems work autonomously and interact with one another, creating machine-centered feedback mechanisms that can cause unexpected outcomes. Semi-autonomous and autonomous vehicles, sensor-enabled heavy machinery, drones, robots, and a number of other devices and equipment will increasingly have AI embedded within them. Human's inability to take control of these semi-autonomous or autonomous systems introduces major control risks [Brundage et al, 2018]. AI systems need to recognise when they are failing and allow a human to take control.

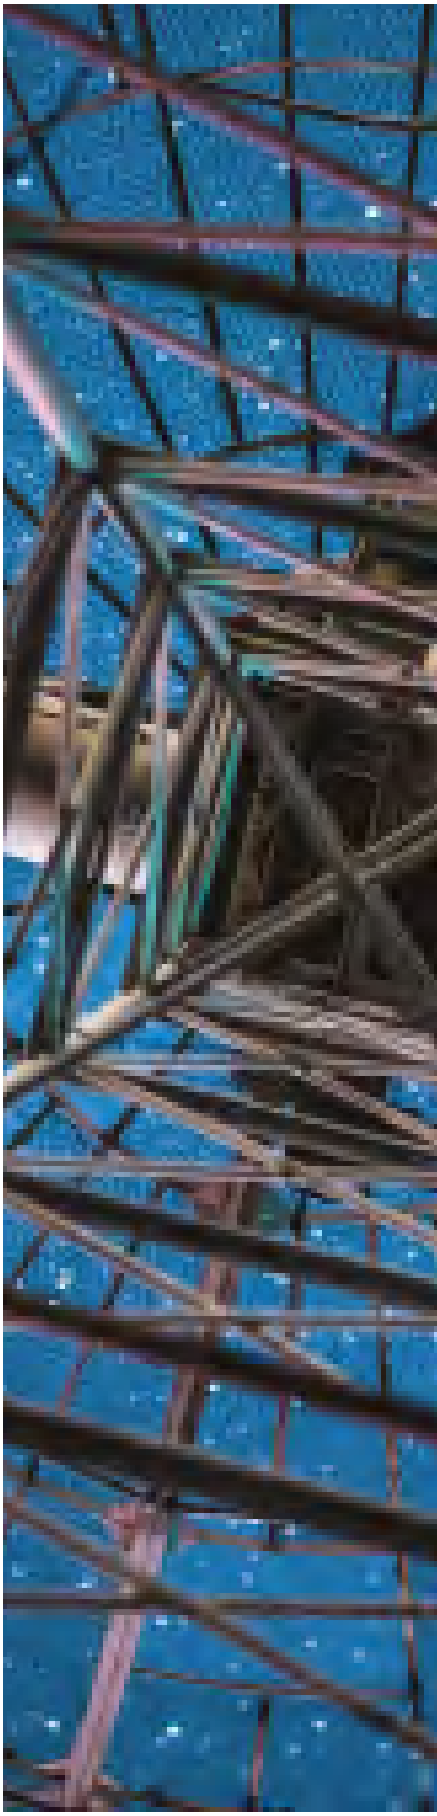

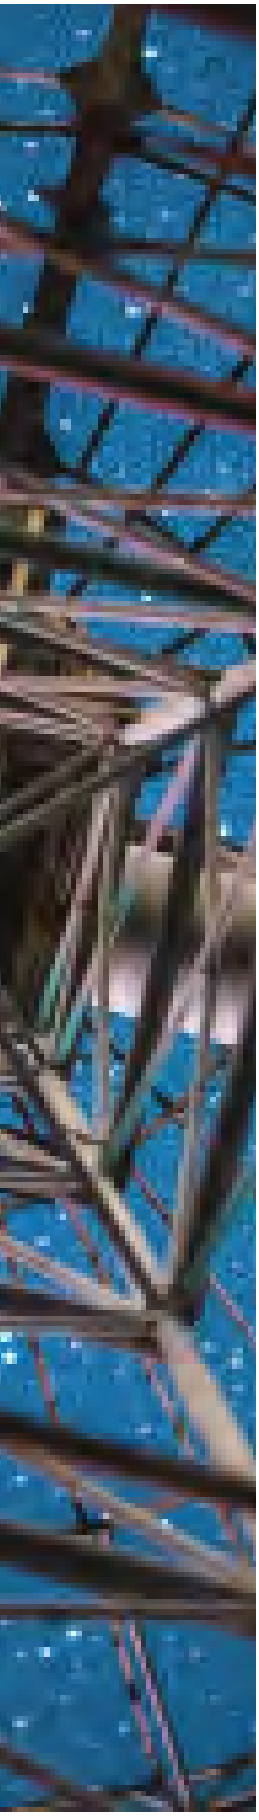

Security risks

Although data security is always a major concern, AI algorithms add a new level of complexity. The more granular the data that is fed to an AI algorithm, the better the algorithm is at personalising a given experience for users. Consumers typically appreciate it when companies can provide personalised experiences tailored to their needs. However, in the process, users' privacy or the confidentiality of their data might be compromised, leading to the need for conscious tradeoffs in security policies.

Misuse of AI via hacking is a serious risk, as many algorithms that are developed with good intentions (for example, for autonomous vehicles) could be repurposed to do harm (for example, turning a self-driving car into a weapon). This introduces new risks for global safety [Brundage et al, 2018] and a call by many AI researchers to ban research on autonomous weapon systems. A large number of AI researchers have signed declarations not to work on such systems.

Good governance is required to build explainability, transparency, and validity into the algorithms [Easterbrook, 2010], including drawing lines between beneficial and harmful AI [Holdren and Smith, 2016]. Machine-learning models — especially the deep learning kind — can also be duped by malicious inputs known as "adversarial attacks." For example, it is possible to find input data combinations that can trigger unwanted outputs from machine-learning models, in effect hacking them. A deep learning model that can recognise stop signs can be duped with additional overlay of data to read the sign as a specific speed limit sign. The harm such hacking could potentially cause is immense, and could lead to an overall lack of trust by consumers in AI systems [Ackerman, 2017]. Not surprisingly, the PwC U.S. 2019 AI Predictions survey [Rao et al, 2019] sound that 64 percent of business executives surveyed rated "boosting AI security" as one of their top prioriies.

Figure 2: AI risks

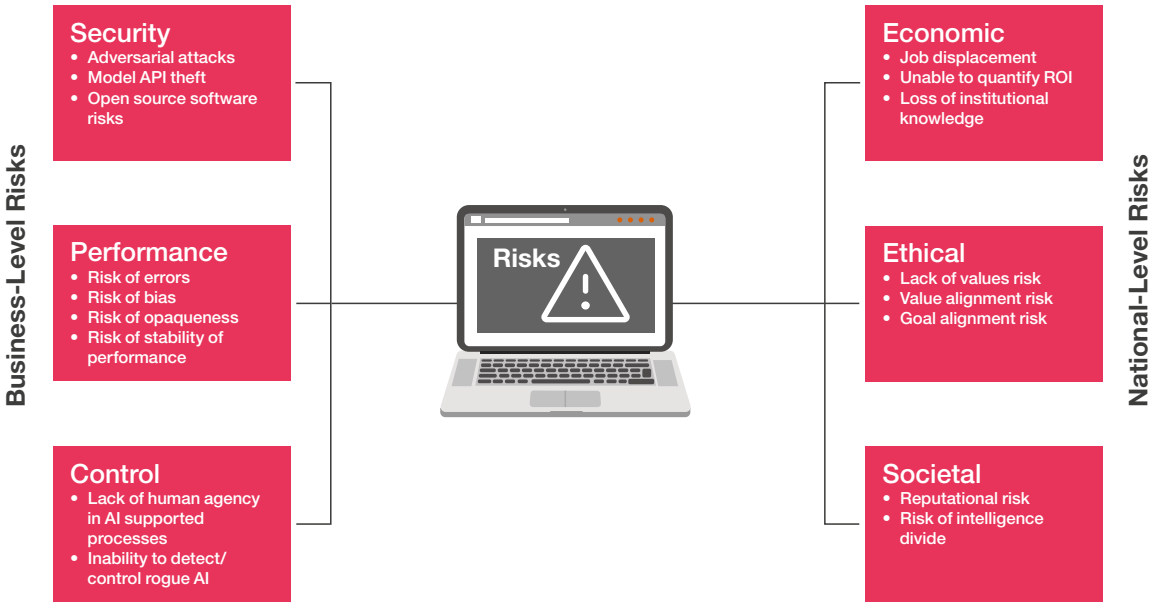

Source: PwC, 2019

3 Managing Trade-Offs

Some countries have started exploring a series of trade-offs that AI presents in an attempt to address them in their policy documents, acknowledging that all of society — businesses, individual consumers, and academics alike — plays a role in how these issues are managed. The trade-offs boil down to three main categories: innovation versus regulation, the individual versus the state, and transparency versus system vulnerability. In all cases, countries — and companies — will have to determine how best to achieve a balance between one side and the other. None of the trade-offs is mutually exclusive, and how to best strike the right balance will depend on a variety of factors.

Innovation vs. regulation.

Regulation can inhibit or accelerate innovation. When regulation clarifies or reduces uncertainties it can give rise to more stability, investment, and innovation. Too many regulations or regulation that is too rigid could stifle companies' ability to introduce new AI applications by, for example, clamping down so much on use of consumer data that they are unable to properly train their algorithms. The complexity is heightened for multinationals operating in different territories with different regulations. The more data available to train an AI system, the smarter the system can become, so territories with less stringent data-use regulations may gain a leg up when it comes to using AI to create custom products or services.

As it relates to AI, regulation can act in both ways. For example, clarifications of liability law in the case of autonomous vehicles and deregulation of the drone airspace could spur more innovation in AI. Conversely, the European Union's GDPR (General Data Protection Regulation) that went into effect in 2018 and other requirements that restrict access to personal data with a view to protecting customer privacys could also inhibit the ability to offer personalisation using machine learning.

Companies can help government officials better understand how much and what type of regulated data they need to properly train AI systems. Companies can also help devise ways to comply with existing consumer protection requirements. Some regulators, such as the U.K.'s Financial Conduct Authority, are experimenting with new approaches — creating a regulatory sandbox, for example. Other countries, such as Canada, are creating AI "superclusters" to attract private funding and retain talent, and to transfer IP from academic labs to commercial enterprises to speed up AI innovation and commercialisation.

The geopolitical dynamics of AI is resulting in three distinct groupings. Some countries are inclined towards more consumer data protection and regulation; while some prefer less regulation regarding the use of consumer data, emphasising common good instead of individual liberty; a third set of countries fall somewhere in between. How countries strike a balance between innovation and regulation will shape how AI advances within each country and region.

The individual vs. the state.

There is a balancing act between individual data privacy, which remains paramount, and the governments' need to access data to further the common good or prevent a malicious act. Still, protecting consumer privacy is a top priority for some governments, which may impact how companies in those nations can use consumer data in their AI systems.

Use of facial recognition and other AI systems that tap into large volumes of personalised data can provide objective security to many citizens. However, in some countries the fear of the loss of liberty far outweighs this objective security. Once again, you might see a divide between some nations that allow the personal use of data for national security, while others will oppose its use.

In the age of social media and smart devices, the volume of available consumer data is massive — and countries will regulate its use in a variety of ways. Some of this will be based on cultural attitudes. In some parts of the world, people are more open to sharing data, while in others there is a greater expectation of privacy protection. PwC's Global CEO Survey found that respondents in Germany, the U.S., and the U.K. are open to regulating how the government collects and uses citizen data, while those in China, India, and Japan favor fewer such limitations.

Transparency vs. system vulnerability.

Government AI strategies may also attempt to balance the need for people to trust AI systems by understanding how they work against the desire to protect the systems from being attacked. The easier it is to explain how the AI "thinks," the logic goes, the easier it becomes for those with malicious intent to infiltrate that system. This will be a major issue in industries such as finance and healthcare, which house massive amounts of sensitive personal data and require a high level of trust between consumers and service providers.

One of the ethical objectives proposed at the 2017 conference on AI principles and widely adopted by businesses is the notion of transparency. For example, a number of institutions are transparent about whether the customer is talking to a human or a chatbot or intelligent assistant (software). While this may be the right thing to do, it also increases the vulnerability of AI systems to malicious attack or adversarial attacks.

Countries need to make a choice about where along the transparency spectrum they want to land. Being too transparent could open them up to malevolent attacks, and they need to be prepared with adequate measures for detecting, defending and prosecuting malicious attacks. However, choosing to be opaque around what decisions are made by AI systems and when they are used could result in a loss of trust with customers, negatively impacting overall adoption.

How the governments of different countries make these trade-offs and how supranational bodies like the World Economic Forum (WEF), IEEE Standards Association, and others influence these decisions will determine if AI will be used for the common good and well-being of all of humanity or if it will be exploited by selected countries and companies to widen the economic disparities within countries and among countries.

4 Call to Action: Focus Areas for National AI Strategies

There are at least 27 countries that have formulated or are in the process of formulating a national AI strategy [Future of Life 2018]. Each country will need to adopt its own approach to a number of aspects of a national AI strategy including, innovation, education and training, regulation and other policies. In addition, it is also essential that there is an international governance framework in place in order to identify and then deal with the issues which cannot be resolved by individual countries. If this framework is absent, there are the downside risks of "weaponised" use of AI (for either political or economic reasons) that can be hugely significant and deeply problematic.

Given the capacity of AI to disrupt so significantly and at such high speed, countries are likely to have to be more specific about the boundaries within which the competitive economy must operate in order to manage the balance between financial success versus sustainable outcomes for society. Put another way, absent a regulatory or governance framework which guides the use of AI by individual actors within the economy, the default will continue to be an emphasis on decision making by business focused on financial results (typically in the short term), and this is very unlikely to be consistent with the overarching societal objectives of the country. This is a significant challenge we already see at play - driving inequality, exclusion and damaging social cohesion. AI has significant potential to accelerate and exacerbate these risks unless more actively managed to achieve sustainable and acceptable outcomes for society.

Given the considerable opportunities and risks associated with AI, national strategies should leverage the opportunities while mitigating the risks. Not addressing the risks and concerns raised in the previous section could lead to a lack of trust in AI systems and eventually inhibit or even kill the current enthusiasm for AI.

At PwC, we have developed an automated natural language processing tool called, National AI Strategies (NAIS) Radar, that analyses all documents on national AI Strategies and performs topic modelling and summarises the key policy recommendations being made by different countries. Figure 3 provides the key topic areas addressed by different national AI strategies, along with details on specific documents and policy areas.

Figure 3: National AI Strategies (NAIS) Radar: Topic Modeling

AI Strategies Map

Govern, Court, Principle, Strategy, Specific, Privacy, **Law**, Determine, Component, Legal, Proposal, Algorithm, Business Service, Engagement, **Activity**, Function, **Management**, Society, **Data** Framework, Expertise, **Governance**, Relevant, **Organisation**, Action, Challenge, Standard, Perspective, Exist, **Individual**, Harm Audit, Objective, Regulatory, Protection, Policy, Procedure, Regulation act Body, Ability, Benefit, Establish context, Manage, Interest, Project

Topic-wise Document Distribution for All

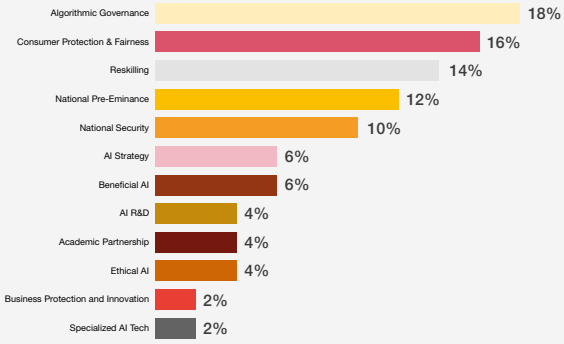

Topic Coverage in ALL for Algorithmic Governance topic

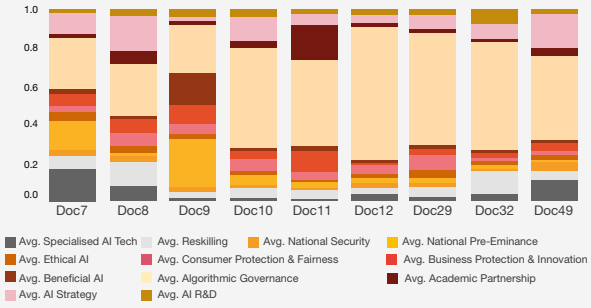

Source: PwC NAIS Model 2019

Innovation

Stimulating AI research and its application to industry is one of the primary goals of many national AI strategies. Countries are doing this using different methods. Some countries are directly funding AI research. For example, the U.K. government has set aside an investment fund of GBP 2.5 billion to be invested in AI and other knowledge-intensive industries [Industrial Strategy 2017]. Others like the U.S., Australia, and others are providing R&D tax credits and other indirect investments to stimulate the research and application of AI technologies in business (See Box 1). Some countries are looking to public-private partnerships to stimulate AI innovation. Earlier this year, the Canadian government launched its Innovation Supercluster Initiative that will invest over $950 million in AI [CAIR, 2017; Canadian government, 2018] (See Box 2). All these efforts are focused on using financial capital to create more cognitive capital. The ultimate expectation is that the cognitive capital will result in more financial capital.

Box 1

Human Rights and AI: Australian Government

The Human Rights Commission launched a major three-year project examining technology and human rights, acknowledging that artificial intelligence, facial recognition, global data markets and other technological developments pose unprecedented challenges to privacy, freedom of expression and equality. An Issues Paper asked how Australian law should protect human rights in the development and use of new technologies, what protections are needed when AI is used in decisions that affect our basic rights and inviting ideas on how we can make technology more inclusive of our diverse community.

In addition to the national initiatives, a number of the state governments have also launched their own initiatives. The Victorian Government has launched an All-Parlimentary AI group to make recommendations. In 2017 New South Wales created a state Data and Analytics centre.

In Queensland, the Artificial Intelligence (AI) Hub is a new initiative under the Queensland Government's landmark $650 million Advance Queensland initiative and is a key part of the Queensland Government's 'Advance Queensland – Foundations for the Future' commitment, focused on the following four key initiatives: (a) build on the existing AI community; (b) develop local AI talent by profiling talented students to employing organisations; (c) provide a launch pad for AI-based startups by connecting them with end-user industries; and (d) assist local industry, corporates and government organisations to better understand and engage with AI technologies.

Source: Matt Kuperholtz, PwC Australia

Box 2

Public-Private Partnerships: Canadian Government

In 2017, the Government of Canada created the Innovation Superclusters Initiative (ISI) to accelerate the growth and development of business-led innovation Superclusters in Canada, translating the strengths of the Canadian innovation ecosystems into new commercial and global opportunities for growth and competitiveness. SCALE AI is one of five successful business-led innovation Superclusters selected to receive federal funding. The combined funding from private sector partners ($700M) and the Federal ($230M) and Provincial ($60M) governments combined is approximately $990M.

SCALE AI's goal is to create highly-skilled jobs and prepare Canada's workforce for the future by skilling and upskilling Canadians to meet an unprecedented demand for digital skills. The development of a next-generation intelligent supply chain will better position Canada as a trading nation and drive productivity and competitiveness of Canadian firms who adopt artificial intelligence.

SCALE AI will protect jobs by providing Canadian firms with a competitive edge in time to market, cost of delivery, supply chain security, and sustainability. SCALE AI will also accelerate the scale up of Canadian businesses, create new businesses, and foster more diversity in the workforce.

SCALE AI is comprised of more than 100 members who are joining forces to apply AI technologies to transform how supply chain is done in Canada. The cluster is meant to bring together academia, start-ups, mid-size, and large businesses with enablers such as technology and advisory firms to work together on joint projects that apply AI in supply chain. This ecosystem is meant to accelerate innovation, create synergies between organisations that would otherwise not collaborate, and encourage innovation. Some examples of applications of artificial intelligence in supply chain include demand forecasting, predictive maintenance, smart scheduling, and automated replenishment and inventory management. The supercluster will drive incremental impact at scale - it is anticipated that more than 16,000 jobs will be created over the next ten years, with an impact to GDP of over 16.5 billion dollars.

Source: Ramy Sedra, PwC Canada

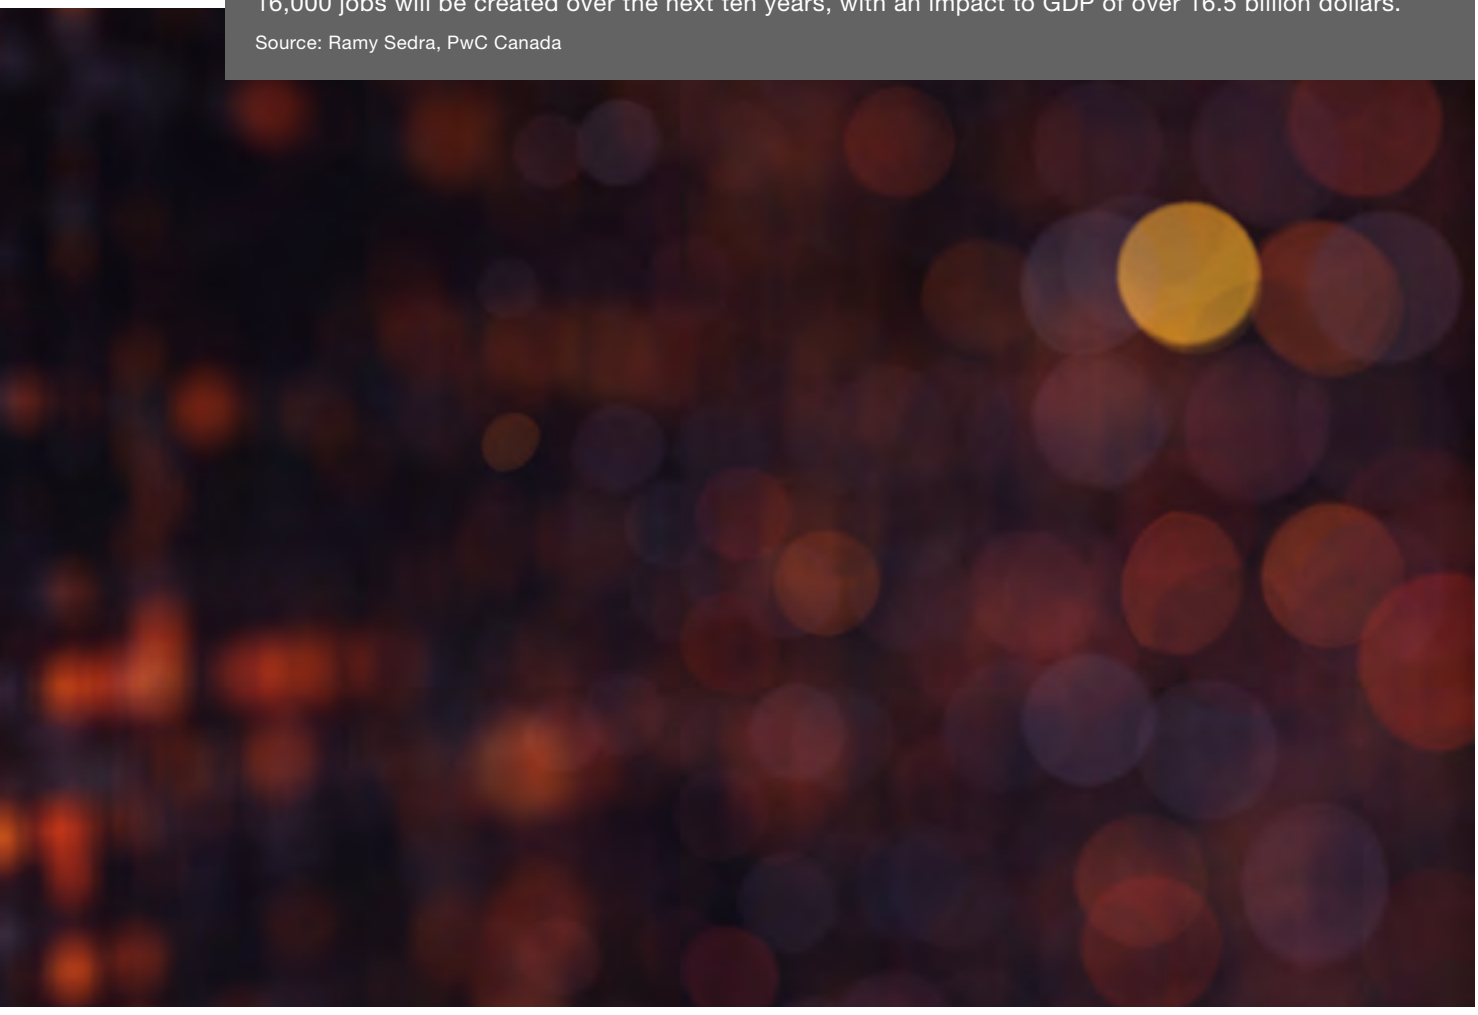

Education & Training

Closely related to innovation is developing the human capital that will fuel the innovation engine and lead to greater behavioural and cognitive capital. Countries are considering three levels of education and training.

Automation may displace 24.7 million jobs by 2027 in the U.S. But the new technologies may also create 14.9 million jobs equating to a total job loss of 17 percent between 2017 and 2027 [Gownder et al, 2017]. As a result, at least one aspect of the training should focus on **reskilling workers** who are displaced due to automation. The European Union is leading the way with a focused effort on reskilling for the fourth industrial revolution [Higgins, 2016; Dittrich, 2016] (See Box 3). In the PwC U.S. 2019 AI Predictions Survey [Rao et al, 2019], 60 percent of executives surveyed rated continual training initiatives for AI as their priority area for reskilling their workforce (See Figure 4).

A number of countries are expanding their AI and data science programmes at the undergraduate, graduate and postgraduate levels. Some universities in the U.S. are introducing undergraduate programmes in AI or developing **specialised AI universities** where AI will be embedded across all disciplines. The U.K. [Hall and Pesentic, 2017] and China [Webster et al, 2017] also have comprehensive programmes to increase AI courses and research in universities (See Box 4).

Emphasis on what children should learn during their primary education in the AI-dominated world of the future is also an active area of interest. The WEF has released a series of policy documents on the "Future of work" and the skills needed.

Figure 4: Building an AI-ready workforce

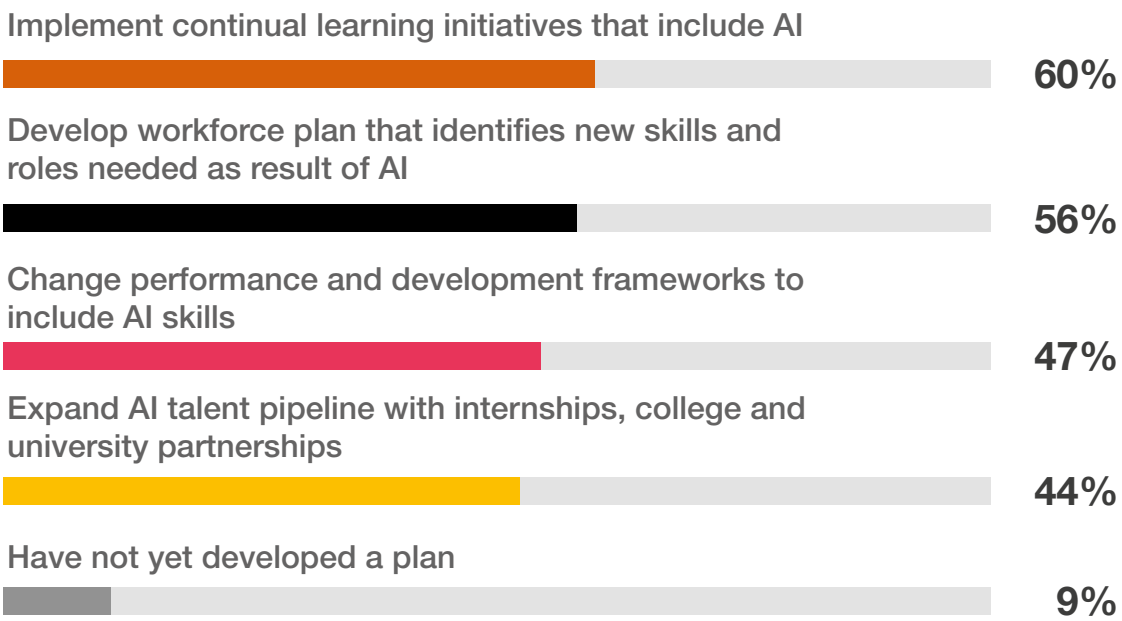

Source: PwC 2019 AI Predictions
Base: 1,001
Q: What are your organisation's plans to adapt to how the nature of work will change as a result of AI in 2019?

Box 3

Reskilling and Upskilling Programmes: Luxembourg Government

In 2018, the Luxembourg government has designed and launched the Luxembourg Skills Bridge pilot programme aimed at accelerating technically and financially the upskilling effort of companies facing a rapid digital transformation. The government's goal is to prevent the consequences of an increasing skills gap for low skilled workers and to boost internal mobility through the accelerated acquisition of "transversal" (problem solving, team work, presentation) skills and digital skills.

The technical government assistance provided to a company is threefold: assistance to (a) better the impact of new technologies on job requirements; (b) assess the present skills portfolio and aspirations of employees (c) identify the best person for the transformed job with the right curriculum to obtain the best matching. The financial assistance covers part of the training costs and consulting costs, but more importantly, the entire salary of the employee during the training period, which could last several months.

More than 15 companies from different sectors and varying in size participated in the programme, illustrating the diversity of digital transformation's impact on blue-collar jobs such as the plumbing trade to highly skilled jobs such as computer science engineers.

The Next Generation Operating Model (NGOM), one of the most sophisticated industry4.0 projects in Europe, led by Husky Injection Molding Systems, a worldwide industrial technology provider to the plastics processing sector was unveiled in February 2019. As this programme has gained momentum, Husky is training a larger NGOM team, which is supported by the Luxembourg Digital Skills Bridge Programme.

Source: Laurent Probst, PwC Luxembourg

Box 4

Office of Artificial Intelligence: U.K. Government

The U.K. has experienced a vibrant and diverse AI discourse over the past two years which has included the launch of the All-Party Parliamentary Group on AI and the publication of the House of Lords AI Review.

Estimates indicate that AI can add an additional 10.3 percent (£232 billion) to the U.K.'s GDP by 2030. To achieve this economic growth, the government has published an ambitious industrial strategy that incorporates AI as a "grand challenge" and has seen a subsequent announcement of a £1 billion AI Sector Deal to include the following priorities:

- Making the U.K. a global center for AI and data-driven innovation by investing in R&D, skills and regulatory innovation
- Supporting sectors to boost productivity through AI and data analytics technologies
- Leading the world in the safe and ethical use of data through a new Centre for Data Ethics and Innovation, and strengthening the U.K.'s cyber security capability
- Helping people develop the skills needed for the jobs of the future through investment in STEM skills and computer science teachers, as well as retraining and researching the impact of automation across sectors

To coordinate this activity, the government opened an Office for Artificial Intelligence. The government has also just announced the launch of 16 Centres for Doctoral Training to train up to 1,000 new PhD students.

The public and private sectors are also working closely to develop the AI and advanced analytics capabilities of U.K. in order to spur innovation and enhance their supervisory activities which resulted in the creation of a 'regulatory sandbox'.

In addition, there has been an array of positive developments in the U.K.'s AI sector. There have been announcements of international collaboration with countries such as France, Canada, the UAE and Singapore amongst others. The U.K. will also be working with the World Economic Forum on an AI procurement policy.

Source: Rob McCargow, PwC U.K.

Regulation

Regulators are concerned with a number of AI risks described earlier. Currently, regulators are in the early stages of understanding the opportunities and risks posed by AI. A number of regulators are actively involving businesses, consumers, academics, and nonprofit bodies in debating a variety of issues, including the following:

Data: Most forms of machine learning require large quantities of data. Business practices tend to vary in terms of the amount of information that is created, how long data is retained, when and with whom it is shared, and what consent is obtained from the users regarding the intent and use of personal data. The GDPR is the most comprehensive regulation targeted at protecting personal data. The principles behind the GDPR enforce transparency, accountability, right to explanation, right to anonymity, purpose limitation, and storage limitation [EU GDRP, 2018].

Fairness: Ensuring that decisions made by AI or automated systems do not discriminate based on gender, age, race, or other factors, where appropriate, is an active area of academic research and public debate. The notion of fairness, is closely tied to the notion of bias. The data used to train machine learning models could be biased. Historical data used to train these models could institutionalise the discrimination in our society, as evidenced by the COMPAS model [Angwin et al, 2016].

Accountability, Transparency, Explainability: A number of issues around who takes responsibility for the performance of an AI system, especially when its recommendation or action is not intuitive to humans are coming to the forefront. The right to an explanation of a machine decision is now a requirement under the GDPR [EU GDPR, 2018]. While any company that has customers in the EU must comply with the GDPR, it is being adopted more widely by other regulators as a best practice model.

Safety, Security and Reliability: All AI systems must be safe, secure and perform reliably throughout their operational lifetime. In addition, human operators should know when an AI system fails (for example, failure transparency) and be in a position to take control. Issues related to failure transparency and human control become extremely important when there is autonomous or partially autonomous vehicles where human drivers are required to be ready to take control, even if the car is on ‘autopilot’.

Surveillance, Privacy and Civil Liberties: Civil liberty groups around the world are becoming increasingly concerned about protecting the privacy of consumers and citizens either from businesses or from government agencies in the name of national security. Organised ways of altering public opinion on critical issues, such as elections, using social media and personalisation engines is also becoming an increasing concern of nations [Scott, 2018].

Policies and standards

Risks associated with performance, control and security may be mitigated by enacting different types of regulations that enforce the governance of data, model, and how AI systems are built, tested and used. However, issues related to economic, societal, and ethical risks may be very difficult to regulate or legislate at this stage. Coming to some level of consensus on certain broad principles and best practices could potentially be a viable approach. Issues that fall under this category may include the following:

Ethics and Values: As the adoption of AI in critical areas of human endeavor increases, the need to align the values of machines with those of humans become critical. Philosophers [Bostrom, 2016] have highlighted the need for value alignment and also the existential risks to humanity if Artificial General Intelligence or Artificial Super Intelligence were to be realised. Various groups and governments are developing principles of conduct around some of them. The Asilomar AI principles lays out 21 principles that address a number of these ethical principles and values [Asilomar AI, 2017]. The IEEE has released principles around ethically aligned design for AI systems [Firth-Butterfield, 2018]. The U.K. government is planning to set aside GBP 9 million towards the establishment of a Center for Data Ethics and Innovation [Industrial Strategy 2017].

Beneficial AI: Increased concerns about the economic impact of AI on jobs and the potential for malicious use of AI has propelled a number of researchers, companies, and organisations to actively look at ways in which AI can be used for the betterment of society, the environment, and disadvantaged citizens [Beneficial AI, 2017].

Autonomous Weapons: The potential for lethal autonomous weapons systems in modern warfare and their impact on human lives have prompted a number of AI researchers to take a proactive stance by refusing to work on such initiatives and actively lobbying for a ban on such systems research. More than 3,100 individuals and 242 organisations have signed the lethal autonomous weapons pledge [Autonomous Weapons, 2018].

Specialised Sectors

While AI will have a transformative impact on all industry sectors, there are some sectors where AI is playing a more disruptive role. Countries that can lead this disruption or that may potentially be impacted by this disruption are taking a more proactive role and creating some related sector-specific initiatives.

Autonomous Vehicles: Autonomous cars and autonomous trucks are AI-fueled disruptions that are expected to fundamentally change the way people and goods move in the future. Countries like the U.S., Germany, the U.K., and Sweden, which all have large auto manufacturing industries, are enacting legislation or clarifying liability issues to foster more innovation and experimentation. Germany released a set of twenty ethical principles for automated and connected driving recently [Ethics Commission, 2017]. There are at least nine bills related to autonomous driving in the U.S., where more than 22 states have enacted legislation regarding the testing and running of self-driving vehicles.

Autonomous Drones: Use of small unmanned flying vehicles, or drones, is another AI-powered sector that could revolutionise delivery of goods. The U.S. Federal Aviation Authority has released rules that will make the use of small unmanned aircraft more feasible [Dorr, 2018]. Deregulation of airspace for drones could lead to significant opportunities for a new drone ecosystem to emerge.

Robotics: Given the expansion of industrial robots and service robots, countries like Japan have released a robot strategy clarifying the vision, use, and deployment of robots, and how they should interact with humans [New Robot Strategy, 2015].

This section highlights some of the key areas that governments looking to formulate a national AI strategy should consider. These areas address both the opportunities of AI and strive to help mitigate the risks associated with AI. In addition, they are aimed at enhancing some of the key forms of capital over the long-term to help build a national competitive advantage.

5 Call to Action: Businesses

As more countries release national AI strategies, businesses should follow these developments closely — and get involved in helping their governments shape policies that will impact the ways that AI and related technologies transform the business landscape. Companies should consider joining policy working groups and jointly advancing AI skills and education, as well as pursuing other efforts that help clarify how to balance their business interests with the greater good. Imagine how governments and businesses could come together and create a more just and prosperous world with more transparent and efficient public administration, more effective and accessible healthcare, more liveable cities, and a more sustainable planet.

Companies around the world are already helping shape national AI strategies in meaningful ways. In Canada, which was one of the first countries to release a national AI strategy, companies are investing heavily in the technology so they'll be able to reap the benefits of policy updates sooner. In the EU, businesses are partnering with governments to upskill, reskill, and reassign workers whose jobs have changed as a result of AI initiatives. European executives are also influencing policy as members of the EU's High-Level Expert Group on Artificial Intelligence, whose charter is to come up with recommendations for policy development that address ethical, legal, social, and economic issues related to AI. In the US and Germany, companies with an interest in the autonomous vehicle market have lobbied effectively for laws that allow them to advance their commercial interests and help to establish public safety. (See Figure 5)

Given the massive opportunities and potential risks associated with AI, companies, global bodies, nonprofit groups, citizens, and policymakers must come together to devise the right strategies that consider the various trade-offs in ways that make sense in their country. Not having a coherent, comprehensive national strategy could put future generations at a competitive disadvantage.

Figure 5: Taking steps toward responsible AI

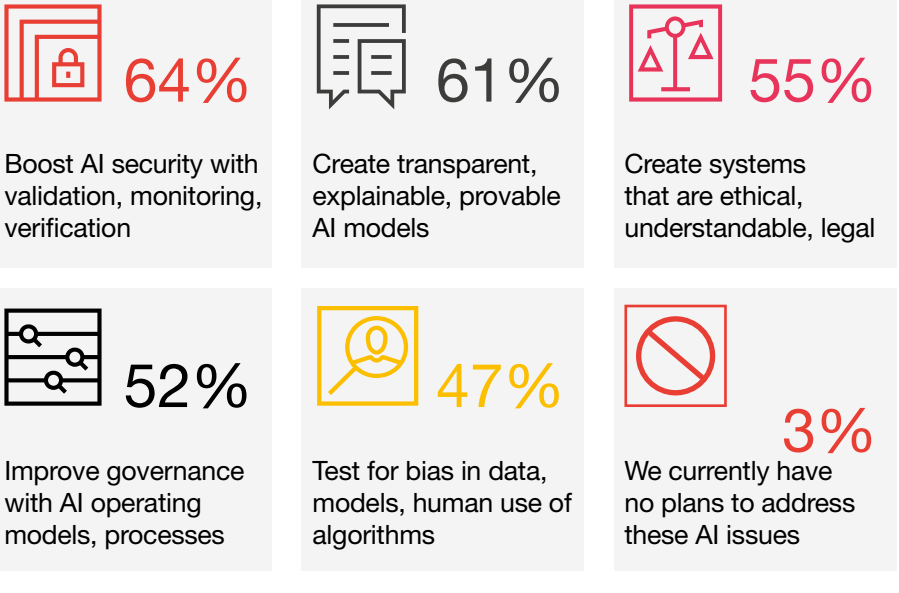

Source: PwC 2019 AI Predictions
Base: 1,001
Q: What are your organisation's take in 2019 to develop and deploy AI systems that are responsible, that is, trustworthy, fair, and stable?

6 References

[Ackerman, E., 2017] Evan Ackerman. Slight Street Sign Modifications Can Completely Fool Machine Learning Algorithms., IEEE Spectrum. August 4, 2017.

[Angwin et al, 2016] Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner. Machine Bias. ProPublica. May 23, 2016.

[Asilomar AI Principles, 2017] Asilomar AI Principles. Future of Life Institute. https://futureoflife.org/ai-principles/. 2017.

[Autonomous Vehicles, 2018] Autonomous Vehicles – Self-Driving Vehicles Enacted Legislation. NCSL – National Conference of State Legislatures. http://www.ncsl.org/research/transportation/autonomous-vehicles-self-driving-vehicles-enacted-legislation.aspx. 2018.

[Autonomous Weapons, 2018] Autonomous Weapons: An Open Letter from AI and Robotic Researchers. Future of Life Institute. https://futureoflife.org/open-letter-autonomous-weapons/. 2018.

[Beneficial AI, 2017] Proceedings of Beneficial AI Conference. Future of Life Institute. https://futureoflife.org/bai-2017/. 2017.

[Bostrom, N. 2014] Nick Bostrom. Superintelligence: Paths, Dangers, Strategies. Oxford, U.K.: Oxford University Press. 2014.

[Brundage et al, 2018] Miles Brundage. The Malicious Use of Artificial Intelligence: Forecasting. Prevention, and Mitigation. University of Oxford. 2018.

[Brynjolfsson, E., and McAfee, A., 2014] Erik Brynjolfsson, and Andrew McAfee. The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies, W. W. Norton & Co. 2014.

[Bughin, J., et al, 2017] Jacques Bughin, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlstrom, Nicolaus Henke, and Monica Trench. Artificial Intelligence – The Next Digital Frontier. McKinsey Global Institute. June 2017.

[CAIR - Canadian Institute for Advanced Research. 2017] Pan-Canadian Artificial Intelligence Strategy Overview, Report, Canadian Institute for Advanced Research (CIFAR), Toronto, Canada, 2017.

[Cave, S., ÓhÉigeartaigh, S. S., 2018] Stephen Cave., and Sean S ÓhÉigeartaigh. An AI Race for Strategic Advantage: Rhetoric and Risks. In Proceedings of AI, Ethics, and Society AIES, 2018.

[Chandler, J., 2009] Jennifer Chandler. Privacy versus National Security – Clarifying the Trade-off. In Kerr, I., Steevers, V., and Lucock, C. Lessons from the Identity Trail. Oxford University Press. 2009.

[Clancey, W. J. 1984] William Clancy. Classification Problem Solving. In Proceedings of the Fourth National Conference on Artificial Intelligence, 49-54. Menlo Park, Calif.: AAAI Press. 1984.

[Dietterich, T. G., 1988] Thomas G. Dietterich. “Artificial Intelligence.” Minsky 1988 2017: 3–24.

[Dittrich, P. J., 2016] Paul Jasper Dittrich. Reskilling for the Fourth Industrial Revolution, Formulating a European Strategy. Policy Paper 175, Jacques Delors Institut, November 3, 2016.

[Dorr, L., 2018] Les Dorr. Fact Sheet – Small Unmanned Aircraft Regulations (Part 107). Federal Aviation Administration. July 23, 2018.

[Easterbrook, S., 2010] Steve Easterbrook. The difference between Verification and Validation. Serendipity, 30 Nov. 2010.

[EU GDPR 2018] Reform of EU Data Protection Rules. European Commission. See https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en. 2018.

[Ethics Commission, 2017] Ethics Commission: Automated and Connected Driving. Federal Ministry of Transport and Digital Infrastructure. June 2017.

[Everson, Miles, Sviokla, J, 2018] Miles Everson and John Sviokla. The Bionic Company. Strategy + Business, February 22, 2019.

[Firth-Butterfield, K. 2018] Kay Firth-Butterfield. Ethically Aligned Design, Technical Report, IEEE, New York City, NY. https://ethicsinaction.ieee.org/. 2018.

[Frey, C.B. and M.A. Osborne 2013] Carl Benedikt Frey and Michael A Osborne. The Future of Employment: How Susceptible are Jobs to Computerisation?' University of Oxford. September 17, 2013.

[Future of Life 2018] Future of Life: Global AI Policy. Future of Life, (https://futureoflife.org/ai-policy/) 2018.

[Gilham, J., et al, 2018] Jonathan Gilham, Lucy Rimmington, Hugh Dance, Gerard Verweij, and Anand Rao. The macroeconomic impact of artificial intelligence. PwC White Paper, PwC. 2018.

[Goodman, B., and Flaxman, S., 2016] Bryce Goodman and Seth Flaxman. EU Regulations on Algorithmic Decision-Making and A 'Right to Explanation. ICML Workshop on Human Interpretability in Machine Learning: 26–30. 2016.

[Govt. of Canada, 2018] Government of Canada's New Innovation program expected to create tens of thousands of middle-class jobs. Innovation, Science and Economic Development, News Release, Feb 18, 2018.

[Gownder et al, 2017] The future of jobs, 2027: Working side by side with robots.

[Gunning, D. 2016] David Gunning. Broad Agency Announcement Explainable Artificial Intelligence (XAI). DARPA. 2016.

[Hall, P., et al, 2017] Patrick Hall, Wen Phan, and SriSatish Ambati. Ideas on interpreting machine learning. O'Reilly Media, 15 Mar. 2017. http://www.oreilly.com/ideas/ideas-on-interpreting-machine-learning

[Hall, D. W., and Pesenti, J., 2017] Dame Wendy Hall and Jerome Pesenti. Growing the Artificial Intelligence Industry in the U.K., Report, HM Government, London, U.K.. 2017.

[Holdren, J., and Smith. M. 2016] John Holdren and Megan Smith. Preparing for the Future in Artificial Intelligence. Executive Office of the President. National Science and Technology Council Committee on Technology. US Federal Government. 2016.

[Industrial Strategy, 2017] Building a Britain Fit for the Future, HM Government. 2017.

[JWYZ, 2017] Graph of Total Number of Websites on the Internet (1991-2015), https://jonathangray.org/2017/10/09/total-websites-on-the-internet/.

[Kelnar, D, 2016] David Kelnar. The fourth industrial revolution: a primer on Artificial Intelligence (AI). Medium, MMC writes, 2 Dec. 2016, medium.com/mmc-writes/the-fourth-industrial-revolution-a-primer-on-artificial-intelligence-ai-ff5e7fffcae1. 2016.

[New Robot Strategy, 2015] New Robot Strategy: Japan's Robot Strategy – Vision, Strategy, Action Plan. Headquarters for Japan's economic revitalization. 2015.

[PwC, 2019] 22nd Annual Global CEO Survey, 2019.

[Rao, A., et al, 2017] Anand Rao, Gerard Verweij, and Euan Cameron. Sizing the Prize: What is the real value of AI for your business and how can you capitalize? PwC White Paper, PwC. 2017.

[Rao, A., et al, 2019] Anand Rao, Scott Likens, Michael Baccala, and Michael Shehab. PwC 2019 AI Predictions: Six AI Priorities you can't afford to ignore. PwC White Paper, PwC, 2019.

[Rao, A. Golbin, I., 2019] Anand Rao and Ilana Golbin. What is fair when it comes to AI Bias? Strategy + Business, April 12, 2019.

[Scott, M., 2018] Mark Scott. Cambridge Analytics helped 'cheat' Brexit vote and U.S. election, claims whistleblower. Politico. March 27, 2018.

[Stewart, L. A., 2010] Luke A Stewart. The Impact of Regulation on Innovation in the United States: A Cross-Industry Literature Review. Institute of Medicine Committee on Patient Safety and Health IT. June 2010.

[Webster et al, 2017] (Translated by) Graham Webster, Paul Triolo, Elsa Kania, and Rogier Creemers, and John Costello.

[Council of the People's Republic of China. 2017] State Council Notice on the Issuance of the Next Generation Artificial Intelligence Development Plan, Report, State Council, Beijing, China. New America, Washington, D.C. 2017.

[WEF 2016] The Future of Jobs – Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution. Global Challenge Insight Report. World Economic Forum. January 2016.

Contacts

United States

Anand S. Rao, PwC Global AI Lead
anand.s.rao@pwc.com

Canada

Ramy Sedra, PwC Canada One Analytics Lead
ramy.sedra@pwc.com

Chile

Jose Triguero, PwC Chile AI Lead
jose.triguero@cl.pwc.com

Brazil

Tomas Roque, PwC Brazil D&A Lead
tomas.roque@pwc.com

Europe

Christian Kirschniak, PwC Europe
D&A Consulting Lead
christian.kirschniak@pwc.com

Luxembourg

Laurent Probst, PwC Luxembourg Economic Development, Digital transformation and Innovation Lead
laurent.probst@lu.pwc.com

Switzerland

Christian Westermann, PwC Switzerland D&A Lead
christian.westermann@pwc.com

Spain

Javier Barguno, PwC Spain D&A Lead
javier.barguno.casanova@pwc.com

Italy

Massimo Pellegrino, PwC Italy AI Lead
massimo.pellegrino@pwc.com

United Kingdom

Euan Cameron, PwC U.K. AI Lead
euan.cameron@pwc.com

Australia

Flavio Palaci, PwC Global DA & AI Lead
flavio.palaci@pwc.com

Matt Kuperholz, PwC Australia AI Lead
matt.kuperholz@pwc.com

China/HK

Yuan Yao, PwC China/HK D&A Lead
yuan.x.yao@cn.pwc.com

Japan

Yann Bonduelle, PwC Japan D&A Lead
yann.y.bonduelle@pwc.com

South East Asia

Scott Albin, PwC SEAC D&A Lead
s.albin@pwc.com

India

Sudipta Ghosh, PwC India D&A Lead
sudipta.ghosh@pwc.com

Africa

Steve Killick, PwC Africa D&A Lead
steve.killick@pwc.com

Middle East

Scott Nowson, PwC Middle East AI Lead
scott.nowson@pwc.com

Author: Dr. Anand S. Rao, Principal, PwC U.S., Global AI Lead

Contributors: Rob McCargow, PwC U.K.; Ramy Sedra, PwC Canada; Laurent Probst, PwC Luxembourg; Matt Kuperholtz, PwC Australia; NAIS Radar team (Ilana Golbin, Vidhi Tembhurnikar, Vishakha Bansal), PwC U.S.

Editor: Leslie Ayers, PwC U.S.; Chi Le, PwC Czech Republic

Layout: Pauline André and Stéphane Denat, PwC Luxembourg

pwc.com/AI

pwc