Artificial Intelligence

A White Paper

Jonathan Kirk, Co-Founder & Director of AI Strategy
systemsGo
www.systemsGo.asia

January 2018

systemsGo

Table of Contents

Big picture

The last thirty years have seen great changes in the world, from the emergence of a globally connected internet, to the pervasive influence of social media. Looking at the advances of the last 10 years, 100 years, and 1000 years, the pace of change is clearly accelerating. But the impact of Artificial Intelligence (AI) has the potential to dwarf the scale of all human progress so far. At the very least, it is likely to bring changes to human society as big or even bigger than mankind's other great transitions, from feudalism to agrarian society, or from a farming culture to industrialized city-dwelling society.

In the last few years, the conversation about AI has reached fever pitch, with some (such as Elon Musk and Stephen Hawking) predicting it will lead to the end of the human race, and others (such as Ray Kurzweil and Mark Zuckerberg) believing that AI could lead to advances in healthcare and other fields that could effectively make humans immortal. Such a wide spectrum of beliefs amongst experts can have the effect

of making it difficult to take any wild claims seriously, and it certainly demonstrates the uncertainty over the impact of AI - are we on the edge of a ‘fourth industrial revolution’, as some claim? Or a precipice that could end mankind? Unquestionably AI has come to dominate the technology discussion, and it is driving huge levels of investment and commitment by investors, companies, and even countries. The news is full of reports that AI will automate more jobs than society can create, or a runaway superintelligence will eliminate humans in favor of a universe of paperclips.

Whatever the ultimate implications of AI, it will bring dramatic changes to the business and social landscape within the next few years. New fields are emerging, new winners are being created, and new rules for success are being written. In this paper, we will try to set out what has changed to breathe life into AI recently, what AI really means, and what the consequences might be. We will show that AI can bring benefits today, and we'll dig into what exactly businesses should be doing today to position themselves to be AI winners.

What is AI anyway?

Part of the challenge of AI is agreeing on a definition. The original intent of the term was the amorphous concept of 'mimicking human intelligence'. As time has gone by, AI has clearly surpassed human levels in a range of different skills (see ‘The state of AI today: better than some humans’ below), and as a result, it is beginning to be usable in situations where previously humans were required.

At the same time, the 'AI' in question is clearly not 'generally' intelligent, and it is not possible to predict today whether it will ever be (see ‘AI to AGI’ below). For mankind, this latter question is of the utmost importance, because it is at the heart of the debate about runaway superintelligence. But it is on a separate track from whether plain 'AI' will continue to affect more and more jobs and situations.

In this paper, by 'AI', we mean a system that can perform one or more specific functions well enough to be comparable with humans, and in general can either train itself or improve its own performance with 'experience' (typically in the form of data).

Why now? What's changed?

The term AI was coined in the 1950s alongside the earliest practical computers, and was born of a grand vision. Despite early results, the field failed to live up to its promise within society's attention span, and the 'AI Winter' of the 1980s stretched out into the

1990s and the new millennium. And then around 2012, something changed, and AI suddenly began to work. And not just work, but work *really well*. Why did this happen? It came from a virtuous cycle of factors: algorithms which began to work once given the right circumstances, the availability of masses of data to train the algorithms, and an explosion of funding, people, and raw computing power working on the field.

Old algorithms, new life

AI systems use a range of algorithms and approaches, but none has caught the imagination or experienced as dramatic an increase in performance as the set of approaches based on neural nets which have come to be called 'Deep Learning'. Neural nets were first described at the outset of the field in the 1950s, and 60s it was proved mathematically that they could potentially solve a variety of problems. In the 1980s a key piece of the puzzle fell into place when Geoff Hinton invented an algorithm that could be used to practically train them - an algorithm known as backpropagation. However, it wasn't until 2012 that another paper by Geoff Hinton showed that with enough data, and enough complexity in the network itself, it could beat the then state-of-the-art image recognition. The field of Deep Learning was born.

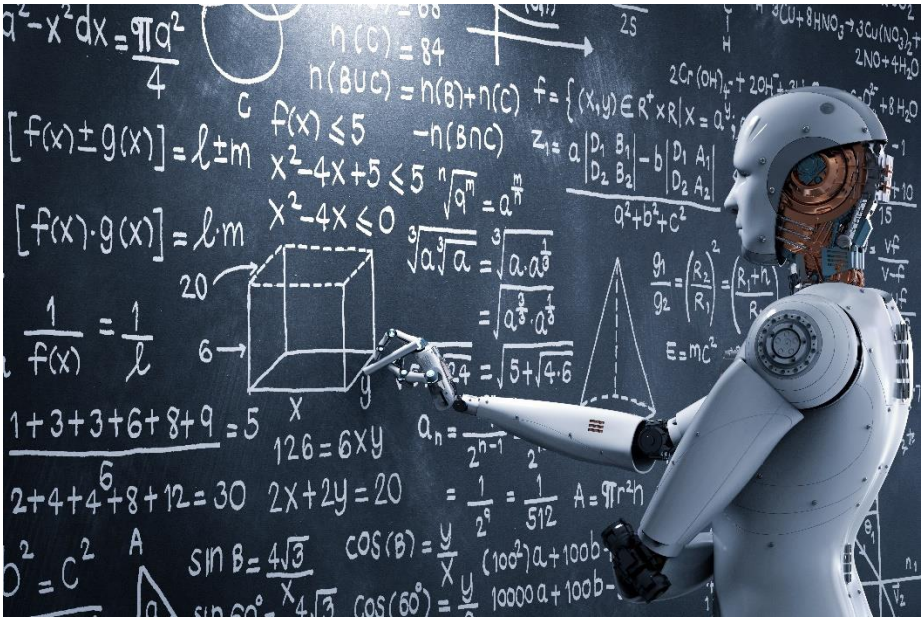

One of the reasons for the excitement about the Deep Learning technique is its (at least superficial) resemblance to the way the brain works, in that there are layers of 'neurons', each of which processes a set of data and distills it to a simpler form, then passes that to successive layers which recognize increasingly abstract patterns. Data is presented as input to the first layer, and then fed through the various layers; the 'output' of the

network, which might be a classification of an image into categories, appears at the output. Nets are 'trained' by presenting them with large volumes of training samples - i.e., samples where the inputs and the desired classification are known. The system learns by processing the training sample, and comparing the result (the desired classification) with the actual result, and 'backpropagating' the error. With enough data, the errors diminish, and the likelihood that the system will accurately classify a *new* input rises - in some cases, past the probability that a human would accurately classify the input.

Neural nets have several flavors, better suited to different tasks. For example:

- Convolutional Neural Nets have become the preferred approach for image recognition.
- Recurrent Neural Networks treat 'time' as an input, so they are good at recognizing patterns spread over time - such as in video.
- Generative Adversarial Networks, invented in 2014, introduce competition between two systems as a way to train both - this technique has proved excellent for *generating* images as opposed to recognizing them.

One of the biggest recent developments of 2017 has been the success of a new technique called 'Reinforcement Learning', which is actually similar to the way humans learn in the real world. The idea is that a system is given a starting point (such as the rules of a game), and an end goal (e.g. 'win'); the system then simply begins testing actions and looking at whether those actions bring it closer to or further from the goal. Reinforcement Learning has begun to bear remarkable fruit: in December, Demis Hassabis, the CEO of DeepMind, announced its AlphaZero system had learnt chess, shogi, and go - and within four hours, had reached a level superior to the world's best players.

DeepMind is moving fast: in February 2016 they announced they had developed the ability to navigate a maze by sight; in March their AlphaGo go-playing system defeated a leading professional player; in May of 2017 AlphaGo handily beat the world's top go player. As Demis put it, this is "akin to a robot being given access to thousands of metal bits and parts, but no knowledge of a combustion engine, then it experiments numerous times with every combination possible until it builds a Ferrari".

One of the most exciting current trends is that these individual networks, each specialized in a key function such as recognizing an image or parsing a written sentence, can now be treated as functional components and assembled like Lego bricks into much larger more complex systems. This has two powerful effects - it broadens the range of functions that AI systems can address dramatically, and it makes it much easier for developers to create new systems by combining existing ones as 'black boxes'.

Data: big and getting bigger

In 2011, the term 'big data' began to explode in popularity {1}. It refers to the enormous wealth of data - facts, figures, images, videos, and web pages - suddenly starting to become available as a result of new authoring, sharing, and storage tools. Initially this spawned whole new industries focused on analytics and visualization. But then a funny thing happened: it turns out that ideas and algorithms which were invented several decades earlier and had largely lain dormant due to poor performance, when applied to these enormous new troves of data, started to yield almost magical results. In short order, the modern field of AI began to emerge from the big data revolution.

And the explosion of data continues at breathtaking speeds. 90% of the world's data was created in the last two years - at a rate of 2.5 quintillion bytes per day {2}: enough to fill the storage capacity of 10 million 256GB iPhone Xs every single day. The nascent 'Internet of Things', the idea that all devices can be connected and share data and control, will multiply the volume of data being produced by thousands of times. As autonomous cars and robot helpers make their way into the real world, they will start to open up vast new streams of data. By 2025, Google estimates individuals will interact with 5000 connected devices per day, more than 20x today - each of which will lead to a new stream of data {3}.

The emergence of AI has only accelerated interest in data. Many initiatives have sprung up to open access to existing data sources, and to find ways to gather new data.

Some key data initiatives

Kaggle – known for its open competitions which have minted some of the highest paid ML researchers. Acquired by Google in 2017-03.	{4}
Harvard Law School Library Innovation Lab CaseLaw Access Project – a project to put all existing caselaw online and digitally accessible	{5}
Waymo, Uber etc have racked up millions of miles of real-world testing of automomous cars. Recent efforts take data from the 20,000 or so unique real-world scenarios and create virtual combinations enabling testing of billions of simulated miles.	{6}
OpenML – the Open Machine Learning project is a collaborative venture for creating open data sets and tools that use them	{7}
UCI Machine Learning Repository – perhaps the oldest and best known repository of open data sets	{8}

Computing power: continuing to explode

Moore's Law, the observation that transistor count (and hence computing power) doubled roughly every two years, held true for about 37 years until 2012; although transistor count growth has slowed, it continues today at about 2.5 years {9}. At the same time that individual device performance is increasing, the number of devices purchased annually, and the installed base, are also growing.

Finally, the suitability of hardware for processing in a brain-like manner has increased dramatically with the discovery that the maths behind the algorithm used to train neural nets benefits greatly from parallel processing, and the 'GPUs' that power gaming consoles and are used in tandem with CPUs in modern devices are ideally suited to this. Combined, these factors are leading to an exponential increase in total global computing capacity.

Computer capacity is not directly comparable with human brain power, but various estimates suggest that total global manufactured computing capacity exceeded the capability of a single brain a few years ago, and is now equivalent to hundreds or thousands of human brains {10}. The point at which a single 'device' (if that concept remains meaningful as network performance and cloud-based computing continue apace) has power equivalent to a human brain is still some decades away, but it is approaching rapidly.

The state of AI today: better than some humans

Games	• Chess – DeepBlue beat world champion Gary Kasparov in 1997 {11} In 2017, DeepMind's system trained itself to beat the world's best players within a few hours {12} • Go – AlphaGo beat the reigning world champion in 2017-05 {13} • Poker - Lengpudashi system beat human players for the first time in late 2017 {14}
Language parsing	• AI matched human parsing capability (around 95%) on sentences of any length in 2016 {15}
Question answering	• As of 2017-11 AI has achieved accuracy of about 78% vs human performance of about 82% {15}
Language sentiment analysis	• OpenAI mLSTM system announced 2017 is able to match human analysis of text sentiment – and is also able to

	generate new text with customizable sentiments {16}
Language translation	• AI achieved a BLEU score of about 0.4, roughly equivalent to human level, in German to English in 2017 {15}
Image recognition	• Character recognition (OCR) • Face recognition • General purpose image recognition
Image generation	• The famous Google X project that learned to recognize (and more importantly, create a representation of) cats {17} • pix2pix and edges2cats – open source systems built on TensorFlow that can turn line drawings into faces and cats
Video generation	• Nvidia system for automatically converting a video of a winter snow scene to a summer scene {18}
Speech recognition	• AI matched human performance of 95% in 2017 {15}
Voice generation	• Google's Tacotron 2 released 2017-12 has near-human accuracy reading from text {19}
Music generation	• MusicGo automatic rap generation by Alibaba {20}
Software generation	• AI Programmer, a fully autonomous software generator, was published in 2017-11 {21}
Theorem proving	• Roughly 80% of state of the art AI theorem solvers can solve all problems in the TPTP (Thousands of Problems for Theorem Provers) dataset

AI to AGI

Although AIs have reached human level performance in many discrete skills, it is readily observed that outside those specific skills, the AIs typically fail completely. AIs lack 'common sense' - the kind of intelligence that humans really mean when we talk about intelligence.

Artificial General Intelligence, or AGI, describes the idea of a more comprehensive human-like intelligence, that enables us to perform many different tasks, and to step back from a problem and 'think' about how to solve it. AGI might (or might not) describe the essence of 'consciousness' - the nature of which is a hard philosophical as well as technical problem that we do not appear to be close to understanding yet (in fact, one of the biggest debates in AI is whether we will need to understand this to develop AGI, or developing AGI will help us to understand it). Whatever its relationship to consciousness, AGI almost certainly requires elements that are not part of the functional AI systems we see today, such as memory, 3D vision, motor control, planning, an understanding of time, a sense of place and navigation. AGI would need the ability to develop and utilize knowledge that does not exist in digitized or easily accessible form - as is still the case with much of human knowledge today. Each of these elements is a

research frontier in its own right, and putting them together is considered the holy grail of AI.

If and when we accomplish that, we would have a machine with the capabilities of humans, but with key advantages - ability to create copies of itself, to work tirelessly, to directly control resources it can dedicate to its own improvement. As a result, there is a very real possibility that 'AGI' could lead very rapidly to 'superintelligence' - something that doesn't exist in human form, and that could change the world order far more fundamentally than through transforming the nature of work.

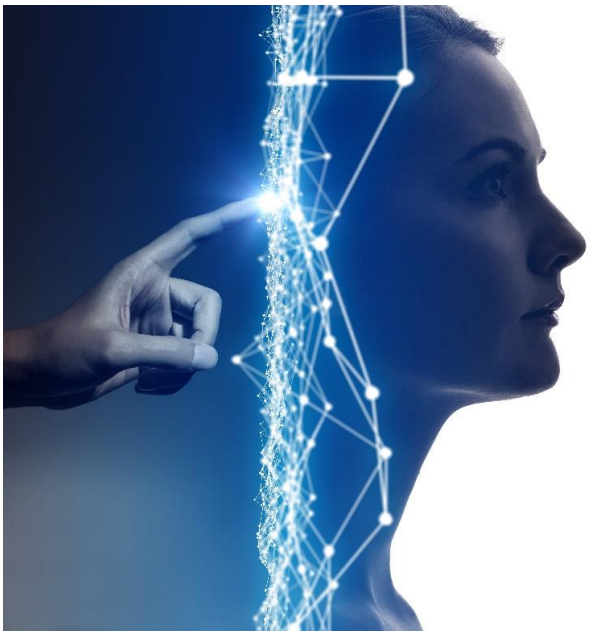

Can this be accomplished? And if it is, will it lead to a world in which humans are no longer necessary, or even a threat? On the nay side, machines clearly learn very differently from humans - we are not exposed to millions of training samples before learning to talk; we really don't actually work very well with 'big data'. Also on the nay side, we actually don't understand very well yet what does drive learning, and we don't understand at all what consciousness is. So we have no real way of knowing whether even the most sophisticated brain emulation will functionally operate in the same way as our own brains. On the other hand, it is possible that 'common sense' is really just an amalgamation of the kinds of individual skills, and that as those skills are acquired - however that it is done - a sufficiently complex system that puts them together might spontaneously develop self-awareness. It is also possible that computers simple keep acquiring skills, and putting them together in ever more complex ways, steadily addressing more varieties of 'but machines can't do that ...' type of challenge, going way past the point of human intelligence in terms of raw capabilities - but still never becoming 'self-aware'.

From a philosophical and policy perspective, this question is very important, because it represents an existential threat to the human race - however low or high the probability of that happening compared to a random asteroid strike. There may only be one chance to get 'superintelligence' wrong, and we need to equip ourselves ahead of time as best

we can to make sure that doesn't happen. "Too many people think the frontiers of AI are delineated by harmless search engines, smart phones, and now Watson. But AGI is much closer to nuclear weapons than to video games." (science fiction author Douglas E. Richards {22}) Or as Stuart Russell, a leading AI researcher, put it “Saying AGI will never happen is like saying we are driving towards the cliff but we’re bound to run out of gas before we get there.” {22}

But AGI is still an uncertain possibility, where the impact of AI techniques on society and business is most definitely not. AI will have a profound and far-reaching impact on the nature of work in the coming decades, independently of what AI ultimately becomes. We will examine this aspect of AI, and its implications for businesses today and in the next few years, in depth.

AI at work

AI today can be thought of as a set of techniques which help to make better decisions. Most people tend to describe them in terms of their functions - image recognition, voice recognition, language parsing, language translation etc. In a broader sense, though, AI techniques like Deep Learning are essentially advanced prediction and classification algorithms, which are used to see patterns in data, and then apply those patterns to take the best action wherever a decision is required. As a result, AI can be used in literally any situation where a decision followed by an action is required. In general business, AI is already being applied in every stage of the product lifecycle, from what to create, how to price, how to market, how to convert, how to distribute, how to support, how to upsell. In almost every industry, AI is making waves in critical applications (see the chart ‘Some examples of AI in use today’).

In the same sort of way that general-purpose tools such as spreadsheets and databases made it dramatically easier to perform calculations and store data, and as a result quickly became pervasive in many different applications, AI will become a part of countless applications used in countless industries.

Some examples of AI in use today - by industry

Automotive / transportation	• Key applications: autonomous vehicles • Waymo is the ‘mindshare’ leader of the self-driving car field, with over 2 million miles driven autonomously, and a robot taxi service expected in 2018 {23} • All major car companies are now investing heavily in autonomous vehicles. Research by Navigant Research ranked the 18 companies most likely to win in autonomous vehicles (in

	order) as: Ford, General Motors, Renault-Nissan, Daimler, Volkswagen, BMW, Waymo, Volvo/Autoliv/Zenuity, Delphi, Hyundai Motor Group, PSA, Tesla, Toyota, ZF, Honda, Uber, nuTonomy, Baidu {24} • Autonomous truck company Otto made a splash when it was bought by Uber for $680m in 2016 only a few months after it was founded (although the nature of the deal is under dispute and may be fraudulent –litigation is still in progress) {25} • ANA just announced $100m investment in AI {26}
Energy	• Key applications: demand forecasting, load balancing • Google's DeepMind using AI to reduce data center costs by 15% {27}
Farming	• Key applications: crop monitoring, water usage forecasting
Finance	• Key applications: fraud detection, credit analayis, financial analysis, automated trading • Mitsui Sumitomo Insurance Company, Limited announced 2017-12 that it would replace 90% of its internal sales department with AI {28}
Healthcare / medicine	• Key applications: health monitoring, predictive diagnostics, medical imaging and analysis • AI beat a panel of 11 pathologists in analyzing breast cancer images {29} • AI at human levels in diagnosing blood infections - just in time because we're running out of trained microbiologists {30}
Law	• Key applications: document discovery and review, brief preparation and validation • JPMorgan COIN performs legal document review in seconds that would have taken humans 360,000 hours {31} • McKinsey estimates that 22% of a lawyer's job and 35% of a paralegal's can be automated {32}
Manufacturing	• Key applications: supply chain optimization, automated defect detection, energy usage forecasting • Andrew Ng (who formed Google Brain and later ran AI for Baidu) recently founded landing.ai based on the belief that manufacturing is the key next frontier for AI {33}
Retail	• Key applications: personalized shopping, customized recommendations
Space	• Key applications: image and signal analysis • NASA is using AI TO discover planets {34}

Emerging applications

Hundreds of startups (over 600, according to the 2017 AI Index {15}) are working on many aspects of AI. These are some of the areas and companies receiving the most attention.

Advertising, Sales, Customer Relationship Management	Appier, Drawbridge, InsideSales.com, Persaido
AI algorithms / bots / computer vision / text analysis	Affectiva, Ayasdi, Blue River Technologies, Captricity, Chronocam, Clarifai, Cognitive Scale, Digital Reasoning Systems, Eloquent Labs, H20.ai, Mindmeld, Mobvoi, Narrative Science, Numenta, Orbital Insight, Vicarious Systems, Zymergen
Automotive Tech	Nuauto, Nutonomy, Zoox
Business Intelligence & Analytics	Bloomreach, Context Relevant, Crowdflower, DataMinr, DataRobot, Logz.io, Rapidminer, Tamr, Trifacta
Cyber security / e-commerce	Bloomreach, Cylance, Darktrace, Sift Science
Fintech	Alphasense, Kensho Technologies
Healthcare	Babylon Health, Benevolent.ai, Icarbonix, Zebra Medical Vision
IoT (Internet of Things)	Verdigris Technologies, Sight Machine
Robotics	Anki, Rokid, Ubtech Robotics

The future of work and jobs

As AI takes on more and more functions, it is obvious that jobs will change, and some jobs will disappear. Based on analysis of the tasks performed in different types of role, several studies have concluded that 30-50% of jobs could be 'affected' or even 'replaced' within the next 20 years (for example see {35} and {36}).

At the same time, new jobs will be created. There are many unmet needs in the world - diseases not cured, food not equitably shared, scientific problems not solved, products with quality that could be improved, new startup ideas not yet developed, hard civilization-wide challenges such as climate change not addressed. As AI helps us to tackle more complex problems, new jobs will emerge related to these needs. As AI takes functions within jobs, those jobs will adapt and take on new dimensions - perhaps a more ‘creative’ dimension; perhaps a more 'human' dimension: consider the example of the Apple Store - instead of automating sales entirely, Apple actually provides an unusually high ratio of retail workers to customers.

Historically, every technical revolution has led worried commentators to believe jobs will be destroyed and to decry the progress, from Ned Ludd and his band of textile machine destroyers, to Keynes in the 1930s, who thought technology would lead to widespread unemployment. In fact, each of these transitions led to a net increase in the opportunities available - but after a period of uncertainty and difficulty while jobs changed, people were transplanted and reskilled. There were certainly winners and losers.

What is different in this time is that AI will touch many, many occupations at once. It may lead to a kind of 'hollowing out' of occupations - very highly skilled jobs will be less affected, and very low skilled jobs will be less affected (low skill jobs remaining today are hard to automate). AI will hit the multitude of jobs 'in the middle' - a vast swathe of occupations. Jobs such as loan analysts, bank tellers, customer service representatives, telemarketers, traders, paralegals and radiologists are all being affected already, and are examples of jobs where a large proportion of the requirement can be automated. The rate at which this will happen is increasing. An emerging challenge is that in many professions, such as law and accounting, junior roles are used for training – but once those junior roles are automated, the cost of training for senior roles rises significantly.

In the short term, the net contribution to jobs appears to be positive – Gartner recently estimated that 2.3 million jobs would be created by 2020 vs 1.8 million lost {37}, in part due to the wealth of new occupations connected with building and applying AI, from machine learning engineers and data scientists needed to develop systems, and a whole range of new jobs such as ‘automation ethicist’ ‘imagery analyst’ ‘automation economist’ and more who are helping to implement AI systems. Even in companies employing significant automation, it is not yet obvious that will lead to a net loss of jobs – for example, in Mitsui Sumitomo Insurance Group’s recent move to automate 90% of

its sales department, affected staff were redeployed to create *new* sales functions, potentially leading to a net increase in requirements {28}.

In the longer run, even if the net result is more jobs in the future, the transition is likely to be very disruptive. As individuals, we need to think about the kinds of skills that are not easily replaceable by AI. As businesses, we need to act now to understand what AI will do to our field, otherwise competitors implementing AI techniques first or faster could gain an insurmountable advantages. Perhaps most importantly of all, governments need to think about how to manage large scale reeducation and reallocation of employment, and how to protect the workers, even while not protecting their jobs, perhaps through new forms of support such as UBI (Universal Basic Income) or existing ones such as universally accessible healthcare.

Robotic Process Automation: RPA

Perhaps the biggest early winner of the AI field is the newly-formed category of software and services called 'Robotic Process Automation'. The name fires the imagination with images of terminators and bladerunners, but the reality is more prosaic: RPA software can mimic the way humans use one or more software systems in sequence. Akin to an evolution of the early software macros that helped users automate repetitive tasks, or the kind of software that powers automated software testing, the software can monitor how a human transacts a given process, and learn to repeat it automatically. The technique has become extremely powerful with the advent of technologies such as fuzzy data extraction and sentiment analysis that deal with the myriad ‘edge’ cases that historically were hard to automate, and because of the fact that (unlike traditional business process automation) it doesn’t require modification of the underlying systems – the RPA system functions as a 'virtual worker', and if processes can be defined clearly enough, the software can handle them fully autonomously.

Although many of the companies pre-date the term and product category, within the last few years the field has grown to include 38 product vendors and more than 50 professional services firms. Within just a few years, Forrester Research estimates there will be over 4 million robotic workers doing administrative, sales, and support tasks. The market is growing dramatically, expected to rise from a few hundred million dollars today to $1bn (Statista) {38} to $3bn by 2021 (Forrester) {39} and $5bn by 2022 (Research and Markets) {40}.

RPA has begun to take off at such an astounding rate in part because of clear results. Success stories abound of companies achieving 50-90% automation of a significant percentage of processes in key functions {41}.

So, is RPA AI?

RPA has become effective in part due to the same techniques that are making AI effective more generally. Macros and rote task automation have been around for decades, but the number of exceptions limited their usefulness. With image recognition, semantic analysis, translation, speech recognition and so on, RPA has enabled automation to pass the point at which it can handle enough exceptions to be truly useful. In this sense, RPA is one of the central functional use cases of AI.

The second side to this question is whether RPA can evolve beyond process automation and into a larger role in the enterprise. For example, consider the distinction between automating handling of a customer service request (RPA today) vs understanding from a customer's request where he/she is in a purchase process and figuring out how to move that forward (RPA tomorrow enhanced by more AI).

Leading RPA vendors

Automation Anywhere	Several platforms including Automation Anywhere Enterprise, Bot Insight, BotFarm, IQ Bot Founded 2006 as Tethys Solutions	https://www.automationanywhere.com/
Blue Prism	Founded 2001	https://www.blueprism.com/
Contextor	Contextor Interactive, Studio and Galaxy platforms	
EdgeVerve Systems	Subsidiary of Infosys AssistEdge platform	https://www.edgeverve.com
IPSoft	Founded 1998 in New York Amelia cognitive agent IPCenter infrastructure management	http://www.ipsoft.com/
Kryon Systems	Founded 2008 Leo platform launched 2012	https://www.kryonsystems.com/
NICE	NICE Robotic Automation platform	https://www.nice.com
Pegasystems	Pega Robotic Automation & Intelligence platform	https://www.pega.com/
UiPath	Founded 2005 in Romania	https://www.uipath.com
WorkFusion	Smart Process Automation platform	https://www.workfusion.com/rpaexpress

Major builders of AI

One key category of AI job creation is in the companies at work creating AI itself. AI is being developed more collaboratively than any other technology in the history of the world - thankfully, many of the leading companies are sharing their technologies. Here we look at leading companies in the field and their AI-focused efforts and accomplishments, as well as some of the key open technology and sharing initiatives.

Alibaba	• Investing $15bn into AI and other technologies through DAMO Academy • Cloud business roughly $1bn per year {42}
Amazon	• Part of Partnership on AI • Leading cloud platform, AWS – roughly $16bn per year {42}
Apple	• Part of Partnership on AI • Open sourced Turi Create ML framework {43} • Siri platform for 'cognitive conversation'
Baidu	• Created DL lab in 2014 • Beat Microsoft to better-than-human speech recognition {14} • Project Apollo autonomous driving platform - now more than 50 member companies
Google	• Google Brain - behind the TensorFlow open AI framework, which is becoming the leading AI software library • Acquired London-based DeepMind • Many big data initiatives, including the research that led to Hadoop, Dremel, Spanner and others • Part of Partnership on AI • Waymo autonomous vehicle pioneer
IBM	• Watson cognitive computing platform which came to fame in 2011 when it beat expert humans at the question and answer game Jeopardy {44} • Part of Partnership on AI
Facebook	• Part of Partnership on AI • Open Compute initiative for open hardware infrastructure design {45} • FBLearner Flow – company-wide internal AI initiative (but not open!) • Internal AML team, applied to various areas – eg face recognition that can find 1 face in 800 million in less than 5 seconds; language translations
Microsoft	• Part of Partnership on AI • Microsoft Azure cloud platform
Nvidia	• Various systems for dynamic video generation {18}

Tencent	• Formed AI lab in 2016 • Enormous scale: WeChat has 900 million daily active users; League of Legends is played by 100 million people every month • Set up AI research lab in Seattle in 2017-05, led by former Microsoft scientist {46}

Global collaboration and open access

The rise of open source software over the last 20 years has contributed tremendously to the emergence of modern AI. In the last few years, AI has begun to benefit greatly from open source. Driven partly by a recognition that the hard problems involved benefit from more (human) brainpower than any organization can employ, and partly from a fear that without open access and transparency, AI technologies could be developed secretly for military applications, several of the leading players have opened access to core technologies and training systems. The result of this is that more and more people have access to the tools to learn AI, more and more are contributing, and the pace of change is accelerating even further.

Open initiatives and organizations

AAAI	• Association for the Advancement of Artificial Intelligence
Allen Institute for AI (AI2)	• Founded by Paul Allen, co-founder of Microsoft • Several projects focused on using AI to address key maths and science problems
Center for Minds Brains and Machines	• Established at MIT
Coursera	• Founders were behind the famous Stanford AI course
DeepMind Lab	• Open source 3D learning environment for agent-based AIs
Future of Life Institute	• Founded by Max Tegmark, author of 'Life 3.0: Being Human in the age of AI' {47} • Developed the Asilomar AI Principles, a set of guidelines for the future development of 'safe AI', signed by 3500 luminaries in the field
Machine Intelligence Research Institute	• Organization investigating the ethics of emerging AI
OpenAI	• Founded by Elon Musk
Turing Institute	• UK national institute for data science • Founded 2015 jointly by four leading universities (Cambridge, Oxford, UCL, Warwick)
Vector Institute	• Founded by Geoff Hinton, the 'father of AI' • Based in Toronto

	• Focused on developing AI talent and technologies
WEF Future of AI and Robotics Council	• Council for studying and influencing the evolution of AI • Recent think piece: 'the US is losing to China in the AI race' {48}

The rise of actual robots

Like AI, physical robots have been a staple of science fiction for more than a century, although the term 'robot' itself was coined in 1920 by Karel Čapek. The reality has seemed far from the fiction, however, with very little visible impact of robots in daily life. Also like AI itself, there are many signs that robots are finally about to reach the level at which they can become pervasive and useful.

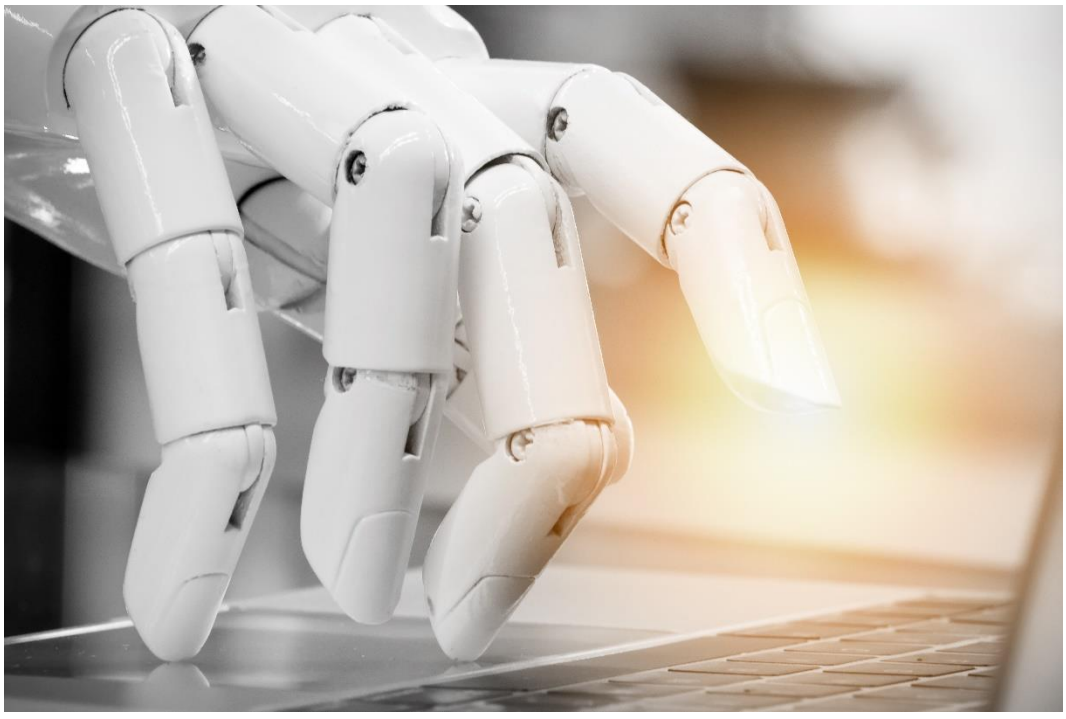

By now most people are aware of the promise of self-driving cars, which are to general purpose robots what a chess-playing AI is to a general purpose AI. After more than a decade of eager anticipation, 2017 saw the first trials of fully autonomous vehicles with no driver, and in Japan, Yamaha demonstrated a riderless motorbike traveling at 200kph {49}. It is possible that within a decade autonomous vehicles will not only

become commonplace but begin to affect public transportation infrastructure and even the design of future urban environments.

One of the key difficulties in creating general purpose robots has been the ability to move around in the same ways and the same sort of environments that humans do. A leader in this field is Boston Dynamics, a company known both for its purchase (and then sale) by Google, and for its regular and slightly creepy videos of robots performing tricks. Towards the end of 2017, BD released a video of its humanoid Atlas robot performing a backflip {50} - in other words, its robots have now arrived at a level of agility few humans can match.

Another key aspect of human interaction with the real world is manual dexterity - our opposable thumbs were a key evolutionary adaptation that enabled humans to begin to develop the technologies that led to the stone age. Recently, a group at Berkeley created a system that uses AI to train a robot to determine the most effective grip for successfully picking up any object of any ship, including objects it has never encountered before. Another key breakthrough in 2017 was the invention of 3D printable synthetic muscles mimicking biological systems, able to lift over 1000x their own weight {51}.

Robots have steadily been making progress into real world applications. Recently Sprint / Softbank announced they will be deploying hundreds of Softbank's Pepper robots to Sprint locations around the US. A Japanese company called Cyberdyne (yes, the same name as the company that invented Skynet in the Terminator movies ...) makes robotic exoskeletons - another example of technologies that augment human activity rather than replacing it entirely {52}.

Also on the Terminator theme, in 2017 the Dubai police announced they had created an actual robotic policeman, which can recognize emotions and hand signals, and separately have developed technologies that can predict the occurrence and location of crime (reminiscent of Minority Report) {53}. One of 2017's most popular viral robotics videos was a short called 'Slaughterbots' which depicted an all-too-believable future defense contractor demonstrating an autonomous killing machine {54} – the video is all the more disturbing because it no longer feels like science fiction. Thankfully, 115 CEOs of major companies joined forces to write an open letter advising against the development of autonomous weapons {55}.

Amazon's acquisition of Kiva for $775m in 2012 and subsequent automation of its warehousing was big news at the time, and by 2017 the renamed Amazon Robotics was operating 45,000 robots and adding over 15,000 per year {56}. Many startups have emerged, and the next frontier of industrial robots are free-moving, vision-guided {57}, and have capabilities far exceeding humans in some dimensions – 7' robot arms from Sarcos can lift 500 lb {58}. In 2017 robots performed cataract surgery for the first time {59}, began to walk on two legs {60}, gained faces almost indistinguishable from

humans {61}, and prepared to compete with humans for jobs in environments historically very hard for robots to operate, such as apple picking {62}

AI in Asia

Asia clearly has the potential to become an epicenter for AI. A recent report by PwC found that China stands to benefit most from the emergence of AI because of the large percentage of its economy that is based on manufacturing. China has announced a massive country-wide investment in AI technology and infrastructure development, with the goal of having AI contribute over $15bn to the Chinese economy by 2020, and being the 'envy of the word' within the next decade {63} {14}. In its comprehensive Global Artificial Intelligence Survey {64}, PwC found that China had more to gain than any other economy from the emergence of AI. Key players in AI have realized the need to involve China in their plans – Google announced in 2017-12 that it would open a new research center in Beijing {65}. The figurative leader of Chinese AI, Kai-Fu Lee, now the head of one of the most visible Chinese venture funds, Sinovation, has over 50 million followers on Weibo {14}.

In Japan, METI (the Ministry of Economy, Trade and Industry) has launched an initiative to rethink the intellectual property and rights framework for AI, to address the contributions and liabilities of stakeholders more fairly given the new considerations AI brings into play {66}. Japan is responsible for some of the biggest innovations, investments and motivations for robotics, and its leadership role looks set to continue with Softbank's clear decision to become a leader in the field, as evidenced by its acquisition in mid 2017 of the US leader Boston Dynamics and the promising young startup Schaft {67}.

For businesses seeking to take advantage of AI, following the progress made by Asian companies is critical. This brings a challenge in dealing with the Japanese and Chinese language environment (although AI is bringing us closer to automatic language translation, it's not at human level yet).

What to do? Key actions for today

Open up data sources

First and foremost: open up access to data. AI feeds on data, and that creates a powerful virtuous circle: the more data, the better you can put AI to work; the more you use AI, the better your products and services; the better your products and services, the

more you can sell; the more you sell, the more data you can gather. A data audit to identify potential sources of data that could lead to valuable applications of AI is the very first step you should take. Much of what we think of as 'data' did not exist a few years ago - the sensors and monitoring systems were not attached, and the necessary storage infrastructure was prohibitively expensive. In your organization, many latent opportunities for data collection most likely exist: find them and unlock them.

Modernize data collection infrastructure

Once the data has been identified, the next step is to modernize (or create) the infrastructure required to access them and make them usable. In some cases this will require new collection sensors or systems. In many cases the data is already available, but not accessible - as recently as 2012, IDC estimated that only 1% of usable data was being accessed and made actionable {68} – sometimes because the current collection mechanism is not accessible for analysis (eg paper); in other cases because there are no sensors or systems in place to gather the data. In other cases it will be a question of piping data streams to data warehouses and utilizing cloud storage to make it manageable - cloud-based systems offer many advantages, and adoption continues to rise (MIT estimates 65% of all data and applications will be cloud based by 2019).

A large proportion of the time spent on any data intensive machine learning application is in preparing the data for use - format conversion, ETL systems (Extract, Transform, Load), filling gaps, cleaning out bad data etc. Careful planning involving awareness of the business context of a data collection exercise is vital for helping to reduce wasted time in this stage.

Modernize and modularize IT

Moving to a modern IT infrastructure based on cloud data storage, data streaming, software as a service, and infrastructure on demand has multiple benefits for AI. Modern IT systems bring benefits in their own right – the latest patches and security updates, ability to offload critical infrastructure management etc. From an AI perspective, vendors are already focused on the opportunities AI brings for their own products as a competitive advantage, and utilizing them brings access to that. Modernizing the underlying systems can help smooth what tends to be most common bottleneck in successful application of AI – massaging the available data into a usable form.

Skills and tasks identification

Although AI can seem almost magical with its uncanny predictive ability, applying AI to business processes is not rocket science: it is a practical issue of identifying and categorizing tasks and processes and the skills required to complete them, and

mapping those to the capabilities of AI. A ‘skills and tasks audit’ is an essential first step in determining precisely what the possible rewards of implementing AI could be, and is something that can be conducted prior to any formal commitment to next steps.

Business goals

Rather than an abstract goal to ‘implement AI’, develop concrete business goals around the practical benefits AI might bring. Benefits might be cost reduction, performance improvement, increased cross-functional communication and awareness, or real-time decision making, for example.

Governance

Data and system governance is a critical challenge as AI systems proliferate, especially in strongly regulated environments such as healthcare or financial services. Data access, analytical expertise, connectivity monitoring, version control, rollback and testing capability are all new challenges that may lead to a skills gap for organizations beginning to address the AI opportunity.

Leadership and expectation management

In any time of great change, team members look to their leaders for encouragement. The changes AI may bring make this especially critical. In a recent survey, 80% of Swedes expressed optimism about the changes robots will bring; another recent study in the US found that 70% of workers expressed worry {69}. One of the key reasons is that in Sweden, workers feel that while *jobs* may change, *workers* will be protected; in the US, the assumption is the opposite. Leaders need to reassure workers that the advent of AI can bring great benefits to companies – which can in turn lead to great benefits to the workers; that it is not simply an issue of replacing or eliminating jobs.

Talent shortage - reference the McKinsey report

AI implementation requires new skills. Some of these skills are highly specialized – the kind of skills required to develop core technologies (jobs in AI are amongst the highest paid in the world at the moment). Many of these skills however are additional skills that workers can acquire to enable them to ‘augment’ their current work – in the same sort of way that most white collar workers can write a presentation or build a spreadsheet. McKinsey estimates that there is already a deficit of 1.5 million suitably skilled managers {70}, and the huge shortage of necessary skills will make it very difficult for companies to fully embrace AI, which will bring enormous benefits to the companies that get there first.

Closing thoughts

Much of the global discussion around AI focuses on the ‘very big picture’ – the potential for AI to change society fundamentally. While this topic is critical for governments, businesses today need to focus on the much more practical reality that AI is going to change the nature of work and competition in the very near future. The best way to address this is to take steps today to take advantage of what AI already offers. Conduct a skills and tasks audit. Review existing infrastructure and opportunities for improvement and data collection. Identify new skills required and take steps to fill any talent gap through retraining or outsourcing. Set concrete goals and take small steps – one at a time. AI can be the source of your next competitive advantage – and an exciting new world of opportunity.

References and further reading

(In the order encountered in the text)

1 Big data - Google trends analysis
https://trends.google.com/trends/explore?date=2009-11-14%202017-12-14&q=%22big%20data%22

2 How Much Data Does The World Generate Every Minute? | IFLScience
http://www.iflscience.com/technology/how-much-data-does-the-world-generate-every-minute/

3 Google report: Your guide to data analytics and machine learning
https://cloudplatformonline.com/g-suite-cio-guide.html

4 UCI Machine Learning Repository
https://archive.ics.uci.edu/ml/index.php

5 JPMorgan Marshals an Army of Developers to Automate High Finance - Bloomberg
https://www.bloomberg.com/news/articles/2017-02-28/jpmorgan-marshals-an-army-of-developers-to-automate-high-finance

6 Waymo's self-driving cars rack up 4 million miles on public roads - The Verge
https://www.theverge.com/2017/11/28/16709104/waymo-self-driving-autonomous-cars-public-roads-milestone

7 OpenML Home
https://www.openml.org/

8 UCI Machine Learning Repository
https://archive.ics.uci.edu/ml/index.php

9 Moore's law - Wikipedia
https://en.wikipedia.org/wiki/Moore%27s_law

10 Global computing capacity – AI Impacts
https://aiimpacts.org/global-computing-capacity/

11 Deep Blue (chess computer) - Wikipedia
https://en.wikipedia.org/wiki/Deep_Blue_(chess_computer

12 DeepMind's AI became a superhuman chess player in a few hours - The Verge
https://www.theverge.com/2017/12/6/16741106/deepmind-ai-chess-alphazero-shogi-go

13 AlphaGo retires from competitive Go after defeating world number one 3-0 - The Verge
https://www.theverge.com/2017/5/27/15704088/alphago-ke-jie-game-3-result-retires-future

14 China's AI Awakening 中国 人工智能 的崛起 - MIT Technology Review
https://www.technologyreview.com/s/609038/chinas-ai-awakening/

15 2017 Annual Report: AI Index
https://aiindex.org/2017-report.pdf

16 OpenAI sets benchmark for sentiment analysis using an efficient mLSTM | TechCrunch
https://techcrunch.com/2017/04/07/openai-sets-benchmark-for-sentiment-analysis-using-an-efficient-mlstm/

17 In a Big Network of Computers, Evidence of Machine Learning - The New York Times
http://www.nytimes.com/2012/06/26/technology/in-a-big-network-of-computers-evidence-of-machine-learning.html

18 Watch an Algorithm Turn Winter Into Summer in Any Video - Motherboard
https://motherboard.vice.com/en_us/article/xwvz9a/watch-an-algorithm-turn-winter-into-summer-in-any-video-image-to-image-translation

19 Google's voice-generating AI is now indistinguishable from humans — Quartz
https://qz.com/1165775/googles-voice-generating-ai-is-now-indistinguishable-from-humans/

20 Alibaba's AI-Powered Shopping Extravaganza – Synced – Medium
https://medium.com/@Synced/11-11-alibabas-ai-powered-shopping-extravaganza-2734b62e7f23

21 AI Programmer: Autonomously Creating Software Programs Using Genetic Algorithms
https://arxiv.org/abs/1709.05703

22 Infinity Born - Douglas E. Richards (taken from the appendix)

23 The companies most likely to get driverless cars on the road first - Business Insider
http://www.businessinsider.com/the-companies-most-likely-to-get-driverless-cars-on-the-road-first-2017-4

24 Navigant Research Leaderboard Report: Automated Driving Navigant Research

https://www.navigantresearch.com/research/navigant-research-leaderboard-report-automated-driving

25 Waymo case reveals Levandowski got $250 million in Uber stock for Otto | TechCrunch
https://techcrunch.com/2017/05/03/waymo-case-reveals-levandowski-got-250-million-in-uber-stock-for-otto/

26 ＡＮＡ、手荷物搬送や案内にＡＩ・ロボ　年100億円投資　　：日本経済新聞
https://www.nikkei.com/article/DGXMZO24638480U7A211C1TI1000/?n_cid=NMAIL007

27 Google uses DeepMind AI to cut data center energy bills - The Verge
https://www.theverge.com/2016/7/21/12246258/google-deepmind-ai-data-center-cooling

28 事務の9割、AIが代替　三井住友海上の営業職（日経)
https://www.nikkei.com/article/DGXMZO25205040281220 17MM8000/

29 These deep learning algorithms outperformed a panel of 11 pathologists | Radiology Business
http://www.radiologybusiness.com/topics/technology-management/these-deep-learning-algorithms-outperformed-panel-11-pathologists

30 Researchers Trained an AI-Equipped Microscope to Diagnose Deadly Blood Infections
https://futurism.com/ai-equipped-microscope-diagnose-deadly-blood-infections/

31 JPMorgan Marshals an Army of Developers to Automate High Finance - Bloomberg
https://www.bloomberg.com/news/articles/2017-02-28/jpmorgan-marshals-an-army-of-developers-to-automate-high-finance

32 Lawyer-Bots Are Shaking Up Jobs - MIT Technology Review
https://www.technologyreview.com/s/609556/lawyer-bots-are-shaking-up-jobs/

33 Andrew Ng’s Landing.ai wants to bring artificial intelligence to the manufacturing industry, starting with Foxconn | TechCrunch
https://techcrunch.com/2017/12/14/andrew-ngs-landing-ai-wants-to-bring-artificial-intelligence-to-the-manufacturing-industry-starting-with-foxconn/

34 NASA's Archived Kepler Data Is Full of Hidden Exoplanets | Big Think
http://bigthink.com/philip-perry/nasa-has-discovered-a-new-exoplanet-using-ai

35 The Future of Jobs - Reports - World Economic Forum
http://reports.weforum.org/future-of-jobs-2016/

36 The Future of Employment - Carl Benedikt Frey & Michael Osborne
https://www.oxfordmartin.ox.ac.uk/downloads/academic/future-of-employment.pdf

37 Gartner Says By 2020, Artificial Intelligence Will Create More Jobs Than It Eliminates
https://www.gartner.com/newsroom/id/3837763

38 Global robotic process automation market 2021 | Statistic
https://www.statista.com/statistics/740440/worldwide-robotic-process-automation-market-size/

39 Worldwide $5 Bn Robotic Process Automation (RPA) Market, 2022
https://globenewswire.com/news-release/2017/09/22/1126150/0/en/Worldwide-5-Bn-Robotic-Process-Automation-RPA-Market-2022.html

40 RPA: Eight Guidelines for Effective Results - Gartner
https://www.pega.com/future-of-work

41 Robotic automation takes off in the Nordics
http://www.computerweekly.com/news/450417014/Robotic-automation-takes-off-in-the-Nordics

42 Alibaba is spending $15 billion on researching quantum computing, AI - The Verge
https://www.theverge.com/2017/10/11/16458486/alibaba-research-investment-fund-15-billion-ai

43 Apple Shares 'Turi Create' Machine Learning Framework on Github - Mac Rumors
https://www.macrumors.com/2017/12/08/apple-turi-create-machine-learning-framework/

44 Watson (computer) - Wikipedia
https://en.wikipedia.org/wiki/Watson_(computer

45 Inside Facebook’s Biggest Artificial Intelligence Project Ever
http://fortune.com/facebook-machine-learning/

46 Baidu, Alibaba and Tencent initiatives will help China ‘aggressively’ narrow AI gap with US | South China Morning Post
http://www.scmp.com/tech/enterprises/article/2118801/baidu-alibaba-and-tencent-initiatives-will-help-china-aggressively

47 The Great AI Paradox - MIT Technology Review - 2017-12
https://www.technologyreview.com/s/609318/the-great-ai-paradox/

48 The US is losing to China in the AI race | World Economic Forum

https://www.weforum.org/agenda/2017/11/the-us-is-losing-to-china-in-the-ai-race

49 Yamaha's Humanoid "Motobot" Rides a Motorcycle at 124 MPH
https://futurism.com/videos/yamahas-motobot-ride-motorcycle/

50 Watch the Boston Dynamics Atlas Robot Do a Backflip. Yes, a Backflip | WIRED
https://www.wired.com/story/atlas-robot-does-backflips-now/

51 Scientists remove one of the final barriers to making lifelike robots | KurzweilAI
http://www.kurzweilai.net/scientists-remove-one-of-the-final-barriers-to-making-lifelike-robots

52 There's a Real Robotics Company Called Cyberdyne And Now I'm Scared
https://paleofuture.gizmodo.com/theres-a-real-robotics-company-called-cyberdyne-and-now-1679979930

53 Dubai Police to enhance smart services with robot police
http://www.computerweekly.com/news/450421866/Dubai-Police-to-enhance-smart-services-with-robot-police

54 'Slaughterbots' film shows potential horrors of killer drones - Nov. 14, 2017
http://money.cnn.com/2017/11/14/technology/autonomous-weapons-ban-ai/index.html

55 AI leaders call for a stop on development of autonomous weapons
http://www.computerweekly.com/news/450424745/AI-leaders-call-for-a-stop-on-development-of-autonomous-weapons

56 Amazon's robot army has grown by 50% - Business Insider
http://www.businessinsider.com/amazons-robot-army-has-grown-by-50-2017-1

57 The technology gap left by Amazon's acquisition of Kiva Systems - The Robot Report
https://www.therobotreport.com/the-technology-gap-left-by-amazons-acquisition-of-kiva-systems/

58 The Guardian GT Is the Most Bonkers Robot on Earth | WIRED
https://www.wired.com/story/the-guardian-gt/

59 For the First Time Ever, a Robot Performed an Operation Inside a Human Eye
https://futurism.com/for-the-first-time-ever-a-robot-performed-an-operation-inside-a-human-eye/

60 Want a Robot to Walk Like You? Don't Expect It to Look Human | WIRED
https://www.wired.com/story/want-a-robot-to-walk-like-you-dont-expect-it-to-look-

human/

61 China showcases worlds most human-like robots in Beijing - Business Insider
http://www.businessinsider.com/china-world-robot-conference-beijing-artificial-intelligence-human-like-2016-10

62 Apple-Picking Robot Prepares to Compete for Farm Jobs - MIT Technology Review
https://www.technologyreview.com/s/604303/apple-picking-robot-prepares-to-compete-for-farm-jobs/?set=604330

63 China's Got a Huge Artificial Intelligence Plan - Bloomberg
https://www.bloomberg.com/news/articles/2017-07-21/china-artificial-intelligence-plan-seeks-59-billion-industry

64 PwC's Global Artificial Intelligence Study: Sizing the prize
https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html

65 Google to Open Beijing AI Center in Latest Expansion in China - Bloomberg
https://www.bloomberg.com/news/articles/2017-12-13/google-to-open-beijing-ai-center-in-latest-expansion-in-china

66 ＡＩの利益、企業どう配分　経産省が指針づくり
https://www.nikkei.com/article/DGXMZO25103530W7A221C1MM8000/

67 SoftBank is buying robotics firms Boston Dynamics and Schaft from Alphabet | TechCrunch
https://techcrunch.com/2017/06/08/softbank-is-buying-robotics-firm-boston-dynamics-and-schaft-from-alphabet/

68 John Gantz and David Reinsel: The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest
https://www.emc.com/collateral/analyst-reports/idc-the-digital-universe-in-2020.pdf

69 The Robots Are Coming, and Sweden Is Fine - The New York Times
https://www.nytimes.com/2017/12/27/business/the-robots-are-coming-and-sweden-is-fine.html

70 Big data: The next frontier for innovation, competition, and productivity | McKinsey & Company
https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/big-data-the-next-frontier-for-innovation

71 Artificial Intelligence and Life in 2030 – Stanford University

https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf

72 McKinsey report on Artificial Intelligence
https://www.forbes.com/sites/louiscolumbus/2017/07/09/mckinseys-state-of-machine-learning-and-ai-2017/#c08a57d75b64

73 AI Will Add $15.7 Trillion to the Global Economy - Bloomberg
https://www.bloomberg.com/news/articles/2017-06-28/ai-seen-adding-15-7-trillion-as-game-changer-for-global-economy

74 PwC's Global Artificial Intelligence Study: Sizing the prize
https://www.pwc.com/gx/en/issues/data-and-analytics/publications/artificial-intelligence-study.html

Contact us

Tokyo

Azabu Amerex Building 3F
3-5-7 Azabudai, Minato-ku Tokyo
106-0041
JAPAN

+81-3-4550-0880

Osaka

Higobashi-Park Building 4F
1-3-3 Kyomachi-bori, Nishi-ku Osaka 550-0003
JAPAN

+81-6-6446-0432

Hong Kong

Unit A&B, 4/F,
Wing Cheong Commercial Building
19-25 Jervois Street, Sheung Wan
HONG KONG

+852 3610 3000

Shanghai

Room 2204-06, Huaihai China Tower,
No.885 Renmin Road, Huangpu District,
Shanghai, 200010
CHINA

+86 21 32535958

Singapore

84A Amoy Street
069903
SINGAPORE

+65-6715-1530

Beijing

Room 2002, Building 2,
First City Mingdundao, External Guangqument
Street, Dongcheng District, Beijing 100035
CHINA

+86-10-8773-0490

Sydney

Suite 2201 Level 22 Tower Two
Westfield 101 Grafton Street
Bondi Junction, NSW 2022
AUSTRALIA

+61-2-8095-6221