



MMC
VENTURES

The State of AI: Divergence

20
19

In partnership with



BARCLAYS

Contents

2 Introduction

6 Summary

Part 1: The Age of AI

16 Chapter 1 – What is AI?

Modern Artificial Intelligence (AI) – “machine learning” – enables software to perform difficult tasks more effectively by learning through training instead of following sets of rules. Deep learning, a subset of machine learning, is delivering breakthrough results in fields including computer vision and language processing.

28 Chapter 2 – Why is AI important?

AI is important because, for the first time, traditionally human capabilities can be undertaken in software inexpensively and at scale. AI can be applied to every sector to enable new possibilities and efficiencies.

40 Chapter 3 – Why has AI come of age?

Specialised hardware, availability of training data, new algorithms and increased investment, among other factors, have enabled an inflection point in AI capability. After seven false dawns since the 1950s, AI technology has come of age.

Part 2: The State of AI

50 Chapter 4 – The race for adoption

AI may be the fastest paradigm shift in technology history. Increasing adoption masks a growing divergence, among nations and within industries, between leaders and laggards.

66 Chapter 5 – The advance of technology

Advances in AI technology are creating new possibilities. Custom silicon is enabling a new generation of AI hardware. Emerging software techniques are delivering breakthroughs in multiple domains and decoupling progress from the constraints of human experience.

82 Chapter 6 – The war for talent

While demand for AI professionals exceeds supply, winners and losers are emerging in the war for talent.

Part 3: The AI Disruptors

96 Chapter 7 – Europe's AI startups

The landscape for entrepreneurs is changing. Europe's 1,600 AI startups are maturing, bringing creative destruction to new industries, and navigating new opportunities and challenges. While the UK is the powerhouse of European AI, Germany and France may extend their influence.

Part 4: Our AI Future

132 Chapter 8 – The implications of AI

AI will have profound implications for companies and societies. AI will reshape sector value chains, enable new business models and accelerate cycles of creative destruction. While offering societies numerous benefits, AI poses risks of job displacement, increased inequality and the erosion of trust.

Introduction

“The future is already here – it’s just unevenly distributed.”

William Gibson

The State of AI 2019: Divergence

As Artificial Intelligence (AI) proliferates, a divide is emerging. Between nations and within industries, winners and losers are emerging in the race for adoption, the war for talent and the competition for value creation.

The landscape for entrepreneurs is also changing. Europe’s ecosystem of 1,600 AI startups is maturing and bringing creative destruction to new industries. While the UK is the powerhouse of European AI, hubs in Germany and France are thriving and may extend their influence in the decade ahead.

As new AI hardware and software make the impossible inevitable, we also face divergent futures. AI offers profound benefits but poses significant risks. Which future will we choose?

Our State of AI report for 2019 empowers entrepreneurs, corporate executives, investors and policy makers. While jargon free, our Report draws on unique data and 400 discussions with ecosystem participants to go beyond the hype and explain the reality of AI today, what is to come and how to take advantage. Every chapter includes actionable recommendations.

Part 1: The Age of AI

We provide an accessible introduction to AI and its applications.

- AI is a way for software to perform difficult tasks more effectively, by learning through practice instead of following rules.
- AI is important because, for the first time, traditionally human capabilities can be undertaken in software efficiently, inexpensively and at scale.
- AI capability has reached an inflection point. After seven false dawns since the 1950s, AI technology has come of age.
- AI has numerous, tangible use cases. We describe 31 across eight industries and highlight why some industries will be affected more than others.



David Kelnar
Report author
Partner & Head of Research
MMC Ventures

Winners and losers are emerging in the race for adoption, the war for talent and the competition for value creation.

For the first time, traditionally human capabilities can be undertaken in software efficiently, inexpensively and at scale.

Part 2: The State of AI

We explain the state of AI adoption, technology and talent in 2019.

- Adoption of AI has tripled in 12 months; AI may be the fastest paradigm shift in technology history. Increasing adoption masks a growing divergence between leaders and laggards. Globally, China leads the race for AI adoption, while sector adoption of AI is uneven and in a state of flux.
- Advances in AI technology are creating new possibilities. Custom silicon is enabling a new generation of AI hardware. Emerging software techniques, including reinforcement learning and transfer learning, are delivering breakthroughs in multiple domains and freeing system design from the constraints of human experience. New, generative AI will reshape media and society.
- Demand for AI talent has doubled in 24 months. Talent, while increasing, remains in short supply with two roles available for every AI professional. Technology and financial service companies are absorbing 60% of AI talent and causing a 'brain drain' from academia. Companies can better attract talent by re-aligning opportunities to professionals' motivations.

Part 3: The AI Disruptors

Drawing on unique analysis, we explore the dynamics of Europe's AI startups.

- Europe is home to 1,600 AI startups. AI entrepreneurship is becoming mainstream, one in 12 new startups put AI at the heart of their value proposition.
- Europe's AI ecosystem is maturing: one in six companies is a growth-stage company. Expect competition, exits and the recycling of capital and talent.
- The UK is the powerhouse of European AI, with a third of the Continent's startups, but Germany and France are flourishing hubs and may extend their influence in the decade ahead.
- Healthcare is a local point for AI entrepreneurship. Activity is thriving given new, transformational opportunities for process automation through AI, and stakeholder engagement.
- Competition for talent, the limited availability of training data, and the difficulty of creating production-ready technology are entrepreneurs' key challenges.

AI may be the fastest paradigm shift in technology history. Increasing adoption masks a growing divergence between leaders and laggards.

The UK is the powerhouse of European AI, but Germany and France may extend their influence.

Continued overleaf...

AI will disrupt business models, require new corporate competencies and change companies' competitive positioning.

Part 4: Our AI Future

We explain AI's profound implications for companies and societies in the decade ahead.

- AI will broaden participation in markets, cause shifts in sector value chains and accelerate cycles of innovation and creative destruction.
- AI will disrupt business models, require new corporate competencies and change companies' competitive positioning.
- AI offers societies significant benefits and risks. AI will transform healthcare, broaden access to goods and services, and increase industrial and agricultural productivity. However, automation may displace jobs while biased AI systems increase inequality. AI will enable the high-tech surveillance state, while autonomous weapons could increase conflict between nations.

Get in touch

AI **MMC Ventures**, AI is a core area of research, conviction and investment. In the last 24 months we've made 20 investments, comprising 50% of the capital we have invested, into many of the UK's most promising AI companies. If you're an early stage AI company, **get in touch to see how we can accelerate your journey.**

www.mmcventures.com

原创力文档

max.book118.com

预览与源文档一致, 下载高清无水印

AI offers societies significant benefits and risks.

AI **MMC**, AI is a core area of investment. Get in touch to see how we can accelerate your journey.



Barclays

Barclays is delighted to partner with MMC Ventures for 'The State of AI 2019'.

Developments in AI are transforming business and society, changing how we live and work. We believe everyone should benefit from these changes. We're committed to ensuring that new technologies are responsibly designed and accessible to all.

AI can enhance our customers' experiences, improve our operational processes and keep our customers safe. Our collaborations with universities, startups and enterprises help us support the latest innovations.

Our AI Frenzy events have united 4,500 AI academics, practitioners, experts and anyone with an interest in AI, while our community of AI Eagles teach colleagues, customers and local communities about AI. Through our national network of Eagle Labs, we offer access to technologies delivering Industry 4.0 – the 'fourth industrial revolution'.

Get in touch if you would like to be part of our AI journey.

Find out more at labs.uk.barclays.ai.



Steven Roberts
Chief Scientific Adviser
Barclays UK Ventures

Summary


Chapter 1: What is AI?

- 'AI' is a general term that refers to hardware or software that exhibit behaviour which appears intelligent.
- Basic AI has existed since the 1950s, via rules-based programs that display rudimentary intelligence in limited contexts. Early forms of AI included 'expert systems' designed to mimic human specialists.
- Rules-based systems are limited. Many real-world challenges, from making medical diagnoses to recognising objects in images, are too complex or subtle to be solved by programs that follow sets of rules written by people.
- Excitement regarding modern AI relates to a set of techniques called machine learning, where advances have been rapid and significant. Machine learning is a sub-set of AI. All machine learning is AI, but not all AI is machine learning.
- There are more than 15 approaches to machine learning. Popular methodologies include random forests, Bayesian networks and support vector machines.
- Deep learning is a subset of machine learning that is delivering breakthrough results in fields including computer vision and language. All deep learning is machine learning, but not all machine learning is deep learning.
- Deep learning emulates the way animals' brains learn subtle tasks – it models the brain, not the world. Networks of artificial neurons process input data to extract features and optimise variables relevant to a problem, with results improving through training.

▶ See page 18 for recommendations

- Machine learning enables programs to learn through training, instead of being programmed with rules. By processing training data, machine learning systems provide results that improve with experience.
 - Machine learning can be applied to a wide variety of prediction and optimisation challenges, from determining the probability of a credit card transaction being fraudulent to predicting when an industrial asset is likely to fail.
- Excitement regarding modern AI relates to a set of techniques called machine learning, where advances have been rapid and significant.**

The Evolution of AI: Deep learning



AI enables human capabilities to be undertaken in software increasingly effectively, efficiently and at low cost.

Chapter 2: Why is AI important?


- AI technology is important because it enables human capabilities – understanding, reasoning, planning, communication and perception – to be undertaken by software increasingly effectively, efficiently and at low cost.
- General analytical tasks, including finding patterns in data, that have been performed by software for many years can also be performed more effectively using AI.
- The automation of these abilities creates new opportunities in most business sectors and consumer applications.
- Significant new products, services and capabilities enabled by AI include autonomous vehicles, automated medical diagnosis, voice input for human-computer interaction, intelligent agents, automated data synthesis and enhanced decision-making.
- AI has numerous, tangible use cases today that are enabling corporate revenue growth and cost savings in existing sectors.
- Applications will be most numerous in sectors in which a large proportion of time is spent collecting and synthesising data; financial services, retail and trade, professional services, manufacturing and healthcare. Applications of AI-powered computer vision will be particularly significant in the transport sector.
- Use cases are proliferating as AI's potential is understood. We describe 31 core use cases across eight sectors: asset management, healthcare, insurance, law & compliance, manufacturing, retail, transport and utilities.
- We illustrate how AI can be applied to multiple processes within a business function (human resources).

► See page 30 for recommendations

AI has numerous, tangible use cases today that are enabling corporate revenue growth and cost savings in existing sectors.

Sector	Core use cases:			
Asset Management	Investment strategy	Portfolio construction	Risk management	Client service
Healthcare	Diagnostics	Drug discovery	Monitoring	
Insurance	Risk assessment	Claims processing	Fraud detection	Customer service
Law & compliance	Case law	Discovery and due diligence	Litigation strategy	Compliance
Manufacturing	Predictive maintenance	Asset performance	Utility optimisation	Churn prediction
Retail	Customer segmentation	Content personalisation	Price optimisation	Control applications
Transport	Autonomous vehicles	Infrastructure optimisation	Fleet management	
Utilities	Supply management	Demand optimisation	Security	Customer experience

Convolutional neural networks are delivering human-level image recognition



Source: <https://www.aig.org/ai>

Chapter 3: Why has AI come of age?

- After seven false dawns since its inception in 1956, AI technology has come of age.
- The capabilities of AI systems have reached a tipping point due to the confluence of seven factors: new algorithms; the availability of training data; specialised hardware; cloud AI services; open source software resources; greater investment; and increased interest.
- Together, these developments have transformed results while slashing the difficulty, time and cost of developing and deploying AI.
- A virtuous cycle has developed. Progress in AI is attracting investment, entrepreneurship and interest. These, in turn, are accelerating progress.


► See page 42 for recommendations

Chapter 4: Adoption

- AI adoption has tripled in 12 months. One in seven large companies has adopted AI; in 24 months, two thirds of large companies will have live AI initiatives. In 2019, AI “crosses the chasm” from early adopters to the early majority.
- AI may be the fastest paradigm shift in technology history. In the course of three years, the proportion of enterprises with AI initiatives will have grown from one in 25 to one in three. Adoption has been enabled by the prior paradigm shift to cloud computing, the availability of plug-and-play AI services from global technology vendors and a thriving ecosystem of AI-led software suppliers.
- Great expectations are fuelling adoption. Executives expect AI to have a greater impact than any other emerging technology, including Blockchain and IoT.
- Increasing overall adoption masks a growing divergence between leaders and laggards. Leaders are extending their advantage by learning faster and increasing investment in AI at a greater pace than laggards.

- Globally, China leads the race for AI adoption. Twice as many enterprises in Asia have adopted AI, compared with companies in North America, due to government engagement, a data advantage and fewer legacy assets.
- Sector adoption is uneven and in a state of flux. 'Early adopters' (financial services and high-tech companies) maintain a lead while 'movers' (retail, healthcare and media) are rapidly catching up. Government agencies, education companies and charities are laggards in AI adoption. Vulnerable members of society may be among the last to benefit from AI.
- AI is advancing across a broad front. Enterprises are using multiple types of AI application, with one in ten enterprises using ten or more. The most popular use cases are chatbots, process automation solutions and fraud analytics. Natural language and computer vision AI underpin many prevalent applications as companies embrace the ability to replicate traditionally human activities in software for the first time.
- Leaders and laggards face different adoption challenges. Laggards are struggling to gain leadership support for AI and to define use cases. Leaders, difficulties, in contrast,
 - ▶ See page 52 for recommendations

One in seven large companies have deployed AI




Source: AI readiness, n = 2,882

What are your organization's plans in terms of artificial intelligence?

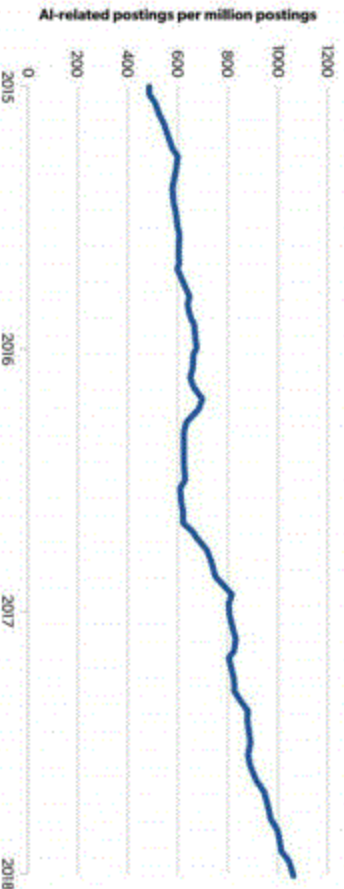
Source: Gartner, 2019 CIO Survey, CIOs have declined to the importance of AI, Figure 1, 2 January 2019

Chapter 5: Technology

- While graphical processing units (GPU) catalysed AI development in the past, and will continue to evolve, hardware innovations are expanding AI's potential. Hardware is being optimised, customised or re-imagined to deliver a new generation of AI accelerators.
- Hardware with 'tensor architectures' is accelerating deep learning AI. Vendors, including NVIDIA and Google are optimising or customising hardware to support the use of popular deep learning frameworks.
- We are entering the post-CPU era. Leading hardware manufacturers are creating new classes of computer processor designed, from inception, for AI. Custom silicon offers transformational performance and greater versatility.
- Custom silicon is also taking AI to the 'edge' of the internet – to IoT devices, sensors and vehicles. New processors engineered for edge computing combine high performance with low power consumption and small size.
- As quantum computing matures, it will create profound opportunities for progress in AI and enable humanity to address previously intractable problems, from personalised medicine to climate change. While nascent,
- Progress in RL is significant because it decouples system improvement from the constraints of human knowledge. RL is well suited to creating agents that perform autonomously in environments for which we lack training data.
- Transfer learning (TL) enables programmers to apply elements learned from previous challenges to related problems. TL can deliver stronger initial performance, more rapid improvement and better long-term results. Interest in TL has grown seven-fold in 24 months and is enabling a new generation of systems with greater adaptability.
- Reinforcement learning (RL) is an iterative approach to developing AI that enables a problem to be solved without knowledge of the domain. Instead of learning from training data, RL systems reward and reinforce progress towards a specified goal. AlphaGo Zero, an RL system developed by DeepMind to play the board game Go, developed unrivalled ability after just 40 days of operation. In 2019, developments in RL will enable groups of agents to interact and collaborate effectively.
- quantum computing is advancing rapidly. Researchers have developed functioning quantum versions of quantum computers.



AI-related job postings as a proportion of total job postings have doubled in 18 months.



Source: Indeed

- By learning fundamental properties of language, TL-powered models are improving the state of the art in language processing – in areas of universal utility. 2018 was a breakthrough year for the application of TL to language processing.

- TL is also enabling the development of complex systems that can interact with the real world, delivering systems with greater adaptability, and supporting progress towards artificial general intelligence, which remains far from possible with current AI technology.

- Generative Adversarial Networks (GANs) will reshape content creation, media and society. An emerging AI software technique, GANs enable the creation of artificial media, including pictures and video, with exceptional fidelity. GANs will deliver transformational benefits in sectors including media and entertainment, while presenting profound challenges to societies – beware ‘fake news 2.0’.

▶ See page 68 for recommendations

Chapter 6: Talent

- Demand for AI talent has doubled in 24 months. There is a gulf between demand and supply, with two roles available for every AI professional.

- The pool of AI talent remains small. AI demands advanced competencies in mathematics, statistics and programming. AI developers are seven times more likely to have a Doctoral degree than other developers.

- Supply is increasing – machine learning has become the top emerging field of employment in the United States. Greater supply is being driven by: high pay; the inclusion of AI modules in university computer science courses; companies’ investment in staff training; and AI technology companies ‘pump priming’ the market with free educational resources.

- Over time, AI tools offering greater abstraction will make AI accessible to less specialised developers.

- Talent shortages are sustaining high salaries. AI professionals are among the best paid developers and their salaries continue to increase; half enjoyed salary growth of 20% or more in the last three years.

- Winners and losers are emerging in the war for talent. The technology and financial services sectors are absorbing 60% of AI talent.

Chapter 7: The European AI landscape

- Europe is home to 1,600 early stage AI software companies. AI entrepreneurship is becoming mainstream. In 2013, one in 50 new startups embraced AI. Today, one in 12 put AI at the heart of their value proposition.

- The 'brain drain' from academia to industry is real and will have mixed implications, catalysing AI's immediate impact while inhibiting teaching and moving value from the public domain to private companies.
- High job satisfaction is intensifying the war for talent. Three quarters of AI professionals are satisfied in their current role.

- To optimise hiring and retention, companies should align roles to AI professionals' primary motivators – learning opportunities, office environment and access to preferred technologies.

- New practitioners in the field are following sub-optimal paths to employment. Company websites and technology job boards are less effective than engaging with recruiters, friends, family and colleagues, according to those already employed in the field.

▶ See page 84 for recommendations



With twice as many AI startups as any other country, the UK is the powerhouse of European AI entrepreneurship

- The UK is the powerhouse of European AI with nearly 500 AI startups – a third of Europe's total and twice as many as any other country. We provide a map of the UK's AI startups and feature 14 leading companies.
- Germany and France are thriving European AI hubs. High-quality talent, increasing investment and a growing roster of breakout AI companies are creating feedback loops of growth and investment.

Availability of talent and access to training data are AI entrepreneurs' key challenges.



The Audio Analytics team listens: Audio Analytic

- Spain's contribution to European AI exceeds its size. Immigration, which correlates with entrepreneurship, has deepened the Country's talent pool.
- The European AI landscape is in flux. While the UK remains the powerhouse of European AI, its share of European AI startups, by volume, has slightly reduced. Brexit could accelerate this. France, Germany and other countries may extend their influence in the decade ahead, spreading the benefits of entrepreneurship more evenly across Europe.
- Italy, Sweden and Germany 'punch above their weight' in core AI technology; while there is support for Nordic countries' reputation for deep tech expertise.
- Nine in ten AI startups address a business function or sector ('vertical'). Just one in ten provides a 'horizontal' AI technology.
- A quarter of new AI startups are consumer companies, as entrepreneurs address or circumvent the 'cold start' data challenge. Many focus on finance or health & wellbeing.
- Healthcare, financial services, retail and media & entertainment are well served by AI startups. In sectors including manufacturing and agriculture, entrepreneurial activity is modest relative to market opportunities.
- Health & wellbeing is a focal point for AI entrepreneurship; more startups focus on the sector than any other. In the coming decade, developers will have a greater impact on the future of healthcare than doctors. Activity is thriving given profound new opportunities for process automation and a tipping point in stakeholders' openness to innovation.
- The UK is the heartland of European healthcare AI, with a third of the Continent's startups. UK entrepreneurs benefit from healthcare scale-ups stimulating talent and increasing openness to innovation within the NHS.
- Marketing and customer service departments enjoy a rich ecosystem of suppliers. A quarter of AI startups serving a business function focus on marketing teams.
- An influx of AI startups supporting operations teams is driving increasing process automation.
- AI companies raise larger amounts of capital, due to technology fundamentals and extensive capital supply.
- Core technology providers attract a disproportionate share of funding. While comprising a tenth of AI startups, they attract a fifth of venture capital.
- AI entrepreneurs' key challenges are the availability of talent, access to training data and the difficulty of creating production-ready technology.

► See page 98 for recommendations

AI offers innovation, efficacy, velocity and scalability

Benefit	Explanation	Examples
Innovation	New products and services.	<ul style="list-style-type: none"> • Autonomous vehicles • Voice-controlled devices
Efficacy	Perform tasks more effectively.	<ul style="list-style-type: none"> • Fraud detection • Customer segmentation
Velocity	Complete tasks more rapidly.	<ul style="list-style-type: none"> • Legal document processing • Manufacturing process optimisation
Scalability	Extend capabilities to additional market participants.	<ul style="list-style-type: none"> • Automated medical diagnosis • Automated executive assistants

Source: MMC Ventures

Chapter 8 : The Implications of AI

- AI's benefits can be abstracted to: innovation (new products and services); efficacy (perform tasks more effectively); velocity (complete tasks more quickly); and scalability (free activity from the constraints of human capacity). These benefits will have profound implications for consumers, companies and societies.
 - By automating capabilities previously delivered by human professionals, AI will reduce the cost and increase the scalability of services, broadening global participation in markets including healthcare and transport.
 - In multiple sectors including insurance, legal services and transport, AI will change where, and the extent to which, profits are available within a value chain.
 - New commercial success factors – including ownership of large, private data-sets and the ability to attract data scientists – will determine a company's success in the age of AI.
 - New platforms, leaders, biggards and disruptors will emerge as the paradigm shift to AI causes shifts in companies' competitive positioning.
-
- AI, 'x-as-a-service' consumption, and subscription payment models will obviate select business models and offer new possibilities in sectors including transport and insurance.
 - As AI gains adoption, the skills that companies seek, and companies' organisational structures, will change.
 - By reducing the time required for process-driven work, AI will accelerate innovation. This will compress cycles of creative destruction, reducing the period of time for which all but select super-competitors maintain value.
 - AI will provide profound benefits to societies, including: improved health; greater manufacturing and agricultural capability; broader access to professional services; more satisfying retail experiences; and greater convenience. AI also presents significant challenges and risks.
 - AI-powered automation may displace jobs. AI will enable the automation of certain occupations that involve routine, in other occupations, AI will augment workers' activities. The short period of time in which select workers may be displaced could prevent those who lose their jobs from being rapidly reabsorbed into the workforce. Social dislocation, with political consequences, may result.

- Biased systems could increase inequality. Data used to train AI systems reflects historic biases, including those of gender and race. Biased AI systems could cause individuals economic loss, loss of opportunity and social stigmatisation.
- Artificial media may undermine trust. New AI techniques enable the creation of lifelike artificial media. While offering benefits, they enable convincing counterfeit videos. Artificial media will make it easy to harass and mislead individuals, and weaken societies by undermining trust.
- AI offers trade-offs between privacy and security. As AI-powered facial recognition advances, to what extent will citizens be willing to sacrifice privacy to detect crime?
- AI enables the high-tech surveillance state, with greater powers for control. China is combining real-time recognition with social scoring to discourage undesirable activity.
- Autonomous weapons may increase conflict. The risk of 'killer robots' turning against their masters may be overstated. Less considered is the risk that conflict between nations may increase if the human costs of war are lower.

► See page 134 for recommendations

"If we fail to make ethical and inclusive AI, we risk losing gains made in civil rights and gender equity under the guise of machine neutrality."

Joy Buolamwini

There are potential harms from algorithmic decision-making


INDIVIDUAL HARMS		COLLECTIVE SOCIAL HARMS
ILLEGAL DISCRIMINATION	UNFAIR PRACTICES	
HIRING		LOSS OF OPPORTUNITY
EMPLOYMENT		
INSURANCE & SOCIAL BENEFITS		
HOUSING		
EDUCATION		ECONOMIC LOSS
CREDIT		
DIFFERENTIAL PRICES OF GOODS		
LOSS OF LIBERTY		
INCREASED SURVEILLANCE		SOCIAL STIGMATISATION
STEREOTYPE REINFORCEMENT		
DIGNITARY HARMS		

Chapter 1

```
class Network(object):  
    def __init__(self, sizes):  
        self.num_layers = len(sizes)  
        self.sizes = sizes  
        self.biases = [np.random.randn(y, 1) for y in sizes[1:]]  
        self.weights = [np.zeros(x, y) for x, y in zip(sizes[:-1], sizes[1:])]  
  
    def feedforward(self, a):  
        for b, w in zip(self.biases, self.weights):  
            a = sigmoid(np.dot(w, a) + b)  
        return a  
  
    def SGD(self, training_data, epochs, mini_batch_size,  
            test_data, n_test = len(test_data)):  
        if test_data: n_test = len(test_data)  
        n = len(training_data)  
        for epoch in range(epochs):  
            random_batches = list(range(0, n, mini_batch_size))  
            mini_batches = [training_data[i:i+mini_batch_size] for i in random_batches]  
            for mini_batch in mini_batches:  
                self.update_mini_batch(mini_batch, n_test)  
            if test_data:  
                print("Epoch (%): (%)/ (%)" % format(  
                    epoch, epoch, n_test))  
                self.evaluate(test_data, n_test)  
            else:  
                print("Epoch (%): complete" % format(  
                    epoch, epoch))  
                print("mini_batch(%i, %i) for %i" % format(  
                    mini_batch_size, n, self.weights))  
        return self.evaluate(test_data, n_test)  
  
    def update_mini_batch(self, mini_batch, n_test):  
        n_b, n_w = len(mini_batch), len(self.biases)  
        delta_n_b = np.zeros(n_b, dtype=float)  
        delta_n_w = np.zeros(n_w, dtype=float)  
        for x, y in mini_batch:  
            delta_n_b, delta_n_w = self.backprop(x, y, n_b, n_w)  
        delta_n_b = np.dot(delta_n_b, n_b)  
        delta_n_w = np.dot(delta_n_w, n_w)  
        self.weights = [w - eta * delta_n_w for w, delta_n_w in zip(self.weights, delta_n_w)]  
        self.biases = [b - eta * delta_n_b for b, delta_n_b in zip(self.biases, delta_n_b)]
```

What is AI?

Modern AI – ‘machine learning’ – enables software to perform difficult tasks more effectively by learning through training instead of following sets of rules. Deep learning, a subset of machine learning, is delivering breakthrough results in fields including computer vision and language processing.



Summary

- AI is a general term that refers to hardware or software that exhibit behaviour which appears intelligent.
- Basic AI has existed since the 1950s, via rules-based programs that display rudimentary intelligence in limited contexts. Early forms of AI included ‘expert systems’ designed to mimic human specialists.
- Rules-based systems are limited. Many real-world challenges, from making medical diagnoses to recognising objects in images, are too complex or subtle to be solved by programs that follow sets of rules written by people.
- Excitement regarding modern AI relates to a set of techniques called machine learning, where advances have been rapid and significant. Machine learning is a sub-set of AI. All machine learning is AI, but not all AI is machine learning.
- Machine learning enables programs to learn through training, instead of being programmed with rules. By processing training data, machine learning systems provide results that improve with experience.
- Machine learning can be applied to a wide variety of prediction and optimisation challenges, from determining the probability of a credit card transaction being fraudulent to predicting when an industrial asset is likely to fail.
- There are more than 15 approaches to machine learning. Popular methodologies include random forests, Bayesian networks and support vector machines.
- Deep learning is a subset of machine learning that is delivering breakthrough results in fields including computer vision and language. All deep learning is machine learning, but not all machine learning is deep learning.
- Deep learning emulates the way animals’ brains learn subtle tasks – it models the brain, not the world. Networks of artificial neurons process input data to extract features and optimise variables relevant to a problem, with results improving through training.

Recommendations

Executives

- Familiarise yourself with the concepts of rules-based software, machine learning and deep learning.
- Explore why AI is important and its many applications (Chapter 2).
- Identify sources of AI expertise, and existing AI projects, within your organisation.

Entrepreneurs

- To identify opportunities for value creation, explore the many applications for AI (Chapter 2) and AI's implications for markets (Chapter 8).
- Familiarise yourself with current developments in AI technology (Chapter 5). New approaches and novel techniques offer new possibilities.

Investors

- Ensure that portfolio company executives are familiar with the concepts of machine learning and deep learning.
- Explore how the limits of rules-based systems are inhibiting portfolio companies. What problems are too complex, or subtle, to be solved by rules-based systems?
- Familiarise yourself with the different approaches to machine learning, to enable you to differentiate between companies deploying meaningful AI and pretenders.

Policy-makers

- AI will impact every industry. Explore Chapter 2 to familiarise yourself with the many applications of AI.
- Explore the positive implications of AI and the risks it poses to society (Chapter 8).

Explore our AI Playbook, a blueprint for developing and deploying AI, at www.mmcventures.com/research.

AI: the science of intelligent programs


Coined in 1956, by Dartmouth Assistant Professor John McCarthy, Artificial Intelligence (AI) is a broad term that refers to hardware or software that exhibit behaviour which appears intelligent. AI is "the science and engineering of making intelligent machines, especially intelligent computer programs" (John McCarthy).

AI is a general term that refers to hardware or software that exhibit behaviour which appears intelligent.

Fig 1. PROSPECTOR Expert System: 1977 Technical Note
(Cover)



Fig 2. PROSPECTOR Expert System: 1977 Technical Note (Detail: Decision Tree)



While expert systems experienced some success, PROSPECTOR predicted the existence of an unknown molybdenum deposit in Washington State; their capabilities were typically limited.

SRI International's PROSECTOR system of 1977 (Fig. 1) was intended to assist geologists' mineral exploration work. Incorporating extensive subject matter information and over 1,000 rules, the system was designed to emulate the process followed by a geologist investigating the potential of a drilling site (Fig. 2).

Basic AI has existed for decades, via rules-based programs that exhibit rudimentary displays of intelligence in specific contexts.

Early AI: rules-based systems

'Expert systems' were a popular form of early AI. Programmers codified into software a body of knowledge regarding a specific field and a set of rules. Together, these components were designed to mimic a human expert's decision-making process.

The limits of rules-based systems

Rules-based systems are limited – because many real-world challenges are too complex, or subtle, to be solved by programs that follow sets of rules written by people. Providing a medical diagnosis, operating a vehicle, optimising the performance of an industrial asset (Fig. 3) and developing an optimised investment portfolio are examples of complex problems. Each involves processing large volumes of data with numerous variables and non-linear relationships between inputs and outputs. It is impractical, and frequently impossible, to write a set of rules – such as a set of ‘if...then’ statements – that will produce useful and consistent results.

Machine learning: software that learns through training

What if the burden of finding solutions to complex problems can be transferred from the programmer to their program? This is the promise of modern AI.

Excitement regarding modern AI relates to a set of techniques called machine learning, where advances have been significant and rapid. Machine learning is a sub-set of AI (Fig. 4). All machine learning is AI, but not all AI is machine learning.

Machine learning shifts much of the burden of writing intelligent software from the programmer to their program, enabling more complex and subtle problems to be solved. Instead of codifying rules for programs to follow, programmers enable programs to learn. Machine learning is the “field of study that gives computers the ability to learn without being explicitly programmed” (Arthur Samuel).


Fig 3. Industrial asset optimisation is a complex problem




Source: Alamy

Machine learning algorithms learn through training. In a simplified example, an algorithm is fed inputs – training data – whose outputs are usually known in advance (‘supervised learning’). The algorithm processes the input data to produce a prediction or recommendation. The difference between the algorithm’s output and the correct output is determined. If the algorithm’s output is incorrect, the processing function in the algorithm changes to improve the accuracy of its predictions. Initially the results of a machine learning algorithm will be poor. However, as larger volumes of training data are provided, a program’s predictions can become highly accurate (Fig. 5).

Fig 4. The Evolution of AI: machine learning



Source: MWC Ventures



Source: Michael Hespan, "Note: The size of data set needed to train a machine learning algorithm is correct (dependent) and correct (generalized)." [https://www.kaggle.com/michaelhespan/note-the-size-of-data-set-needed-to-train-a-machine-learning-algorithm-is-correct-dependent-and-correct-generalized](#)

The defining characteristic of a machine learning algorithm, therefore, is that the quality of its predictions improves with experience. Typically, the more relevant data provided to a machine learning system, the more effective its predictions (up to a point).

By learning through practice, instead of following sets of rules, machine learning systems deliver better solutions than rules-based systems to numerous prediction and optimisation challenges.

There are many approaches to machine learning

There are more than 15 approaches to machine learning. Each uses a different form of algorithmic architecture to optimise predictions based on input data.

- One, deep learning, is delivering breakthrough results in new domains. We explain deep learning below. Others receive less attention – but are widely used given their utility and applicability to a broad range of use cases. Popular machine learning algorithms beyond deep learning include:
 - Random forests that create multitudes of decision trees to optimise predictions. Random forests are used by nearly half of data scientists (Kaggle).
 - Bayesian networks that use probabilistic approaches to analyse variables and the relationships between them. One third of data scientists use Bayesian networks (Kaggle).

- Support vector machines that are fed categorised examples and create models to assign new inputs to one of the categories. A quarter of data scientists employ support vector machines (Kaggle).

Each approach offers advantages and disadvantages. Frequently, combinations are used (an 'ensemble' method). In practice, developers frequently experiment to determine what is effective.

Machine learning can be applied to a wide variety of prediction and optimisation challenges. Examples include: assessing whether a credit card transaction is fraudulent; identifying products a person is likely to buy given their prior purchases; and predicting when an industrial asset is likely to experience mechanical failure.

The defining characteristic of a machine learning algorithm is that the quality of its predictions improves with experience.


Deep learning: offloading feature specification

Even with the power of general machine learning, it is difficult to develop programs that perform certain tasks well – such as understanding speech or recognising objects in images.

In these cases, programmers cannot specify the features in the input data to optimise. For example, it is difficult to write a program that identifies images of dogs. Dogs vary significantly in their visual appearance. These variations are too broad to be described by a set of rules that will consistently enable correct classification (Fig. 6). Even if an exhaustive set of rules could be created, the approach would not be scalable; a new set of rules would be required for every type of object we wished to classify.


Deep learning is delivering breakthrough results in these use cases. Deep learning is a sub-set of machine learning and one of many approaches to it (Fig. 7). All deep learning is machine learning, but not all machine learning is deep learning.

Fig 6. Identifying features can be difficult ('Dalmatians or ice cream?')



Source: Google images


Fig 7. The Evolution of AI: deep learning



Source: MMC Ventures

Even with the power of general machine learning, it is difficult to develop programs that perform certain tasks well – such as understanding speech or recognising objects in images.

Fig 8. Deep learning offloads the burden of feature extraction from a programmer to her program



Source: M&C Ventures


To undertake deep learning, developers create artificial neurons – software-based calculators that approximate, crudely, the function of neurons in a brain.

Deep learning is valuable because it transfers an additional burden – the process of feature extraction – from the programmer to their program (Fig. 8).

Humans learn to complete subtle tasks, such as recognising objects and understanding speech, not by following rules but through practice and feedback. As children, individuals, experience the world (see a dog), make a prediction ('dog') and receive feedback. Humans learn through training. Deep learning works by recreating the mechanism of the brain (Fig. 9) in software (Fig. 10). With deep learning we model the brain, not the world.

To undertake deep learning, developers create artificial neurons – software-based calculators that approximate, crudely, the function of neurons in a brain. Artificial neurons are connected together to form a neural network. The network receives an input (such as a picture of a dog), extracts features and offers a determination. If the output of the neural network is incorrect, the connections between the neurons adjust to alter its future predictions. Initially the network's predictions will frequently be incorrect. However, as the network is fed many examples (potentially, millions) in a domain, the connections between neurons become finely tuned.

Fig 9. A biological neural network



Source: iStock

Deep learning has unlocked significant new capabilities, particularly in the domains of vision and language.

Deep learning enables:


- autonomous vehicles to recognise entities and features in the world around them (Fig. 11);
- software to identify tumours in medical images;
- Apple and Google to offer voice recognition systems in their smartphones;
- voice-controlled devices, such as the Amazon Echo;
- real-time language translation (Fig. 12);
- sentiment analysis of text;
- and more.

Deep learning is not suited to every problem. Typically, deep learning requires large data sets for training. Training and operating a neural network also demand extensive processing power. Further, it can also be difficult to identify how a neural network developed a specific prediction – a challenge of ‘explainability’.

However, by freeing programmers from the burden of feature extraction, deep learning has delivered effective prediction engines for a range of important use cases and is a powerful tool in the AI developer’s arsenal.

Deep learning has delivered effective prediction engines for a range of important use cases and is a powerful tool in the AI developer’s arsenal.

Fig 10. An artificial neural network



Source: NNMC Ventures

Fig 11. Deep learning enables autonomous vehicles to identify objects around them



Source: Museum of Computer Science, MIT, CA

Fig 12. Google’s Pixel Buds use deep learning to provide real-time language translation



Source: Google / Pixel Buds

How does deep learning work?

An artificial neural network is created when artificial neurons are connected together. The output of one neuron becomes an input for another.

Deep learning involves creating artificial neural networks – software-based calculators (artificial neurons) that are connected to one another.

An artificial neuron (Fig. 13) has one or more inputs. The neuron performs a mathematical function on its inputs to deliver an output. The output will depend on the weights given to each input, and the configuration of the input-output function in the neuron. The input-output function can vary. An artificial neuron may be a:


- linear unit (the output is proportional to the total weighted input);
- threshold unit (the output is set to one of two levels, depending on whether the total input is above a specified value);
- sigmoid unit (the output varies continuously, but not linearly as the input changes).

An artificial neural network (Fig. 14) is created when artificial neurons are connected to each other. The output of one neuron becomes an input for another.

Fig 13. An artificial neuron




Fig 14. An artificial neural network



How does deep learning work?

Neural networks are organised into multiple layers of neurons (Fig. 15) – hence ‘deep’ learning. An input layer receives information to be processed, such as a set of pictures. An output layer delivers results. Between the input and output layers are layers referred to as ‘hidden layers’ where features are detected. Typically, the outputs of neurons on one level of a network all serve as inputs to each neuron in the next layer.

Fig 15: Deep learning: structuring an artificial neural network



Source: IBM, Venturus

Neural networks are organised into multiple layers of neurons – hence ‘deep’ learning. An input layer receives information to be processed, such as a set of pictures. An output layer delivers results.

How does deep learning work?


Fig. 16 illustrates a neural network designed to recognise pictures of human faces. When pictures are fed into the neural network, the first hidden layers identify patterns of local contrast (low-level features such as edges). As images traverse the hidden layers, progressively higher-level features are identified. Based on its training, at its output layer the neural network will deliver a probability that the picture is of a human face.

Typically, neural networks are trained by exposing them to a large number of labelled examples. As errors are detected, the weightings of the connections between neurons adjust to offer improved results. When the optimisation process has been repeated extensively, the system is deployed to assess unlabeled images.

The structure and operation of the neural network below is simple (and simplified), but structures vary and most are more complex. Architectural variations include: connecting neurons on the same layer, varying the number of neurons per layer, and connecting neurons' outputs into previous layers in the network ('recursive neural networks').


It takes considerable skill to design and improve a neural network. AI professionals undertake multiple steps including: structuring the network for a particular application; providing suitable training data; adjusting the structure of the network according to progress; and combining multiple approaches to optimise results.

Fig 16. Deep learning: the process of feature extraction




Source: MDC Ventures, Andrew Ng

Chapter 2



Why is AI important?

AI is important because, for the first time, traditionally human capabilities can be undertaken in software inexpensively and at scale. AI can be applied to every sector to enable new possibilities and efficiencies.



Summary

- AI technology is important because it enables human capabilities – understanding, reasoning, planning, communication and perception – to be undertaken by software increasingly effectively, efficiently and at low cost.
- General analytical tasks, including finding patterns in data, that have been performed by software for many years can also be performed more effectively using AI.
- The automation of these abilities creates new opportunities in most business sectors and consumer applications.
- Significant new products, services and capabilities enabled by AI include autonomous vehicles, automated medical diagnosis, voice input for human-computer interaction, intelligent agents, automated data synthesis and enhanced decision-making.
- AI has numerous, tangible use cases today that are enabling corporate revenue growth and cost savings in existing sectors.
- Applications will be most numerous in sectors in which a large proportion of time is spent collecting and synthesising data; financial services, retail and trade, professional services, manufacturing and healthcare. Applications of AI-powered computer vision will be particularly significant in the transport sector.
- Use cases are proliferating as AI's potential is understood. We describe 31 core use cases across eight sectors: asset management, healthcare, insurance, law & compliance, manufacturing, retail, transport and utilities.
- We illustrate how AI can be applied to multiple processes within a business function (human resources).

Recommendations

Executives

- Explore the new possibilities enabled by AI, from voice control and intelligent agents to autonomous vehicles and automated diagnosis, to appreciate AI's importance in the decade ahead.
- Examine AI use cases in a range of sectors to familiarise yourself with the technical capabilities of AI
 - from incorporating additional data into analyses to understanding written and spoken language.
- Identify core aspects of your company's value proposition – for example, analysis or communication
 - to which AI could be relevant.

Entrepreneurs

- AI presents new opportunities to disrupt sectors ranging from manufacturing to healthcare. Identify business processes ripe for improvement or reinvention through AI, particularly in sectors in which data synthesis or processing are extensive.
- AI offers numerous capabilities, from multi-variate analysis to natural language processing. Identify opportunities to use multiple aspects of AI, where appropriate, both within your company and for buyers.

Investors

- Evaluate opportunities and threats to existing portfolio companies from the many applications of AI.
- Seek companies that are using AI to fulfil new possibilities. The paradigm shift to AI will create valuable new winners.
- With AI poised to impact every sector, develop a framework to identify preferred areas of focus and success factors for AI investments.

Policy-makers

- AI presents profound new possibilities, from autonomous vehicles to automated diagnosis. Familiarise yourself with the many applications of AI to appreciate the opportunities it presents and the challenge to regulators it poses.
- Read Chapter 6 of our AI Playbook (www.mmventures.com/research) to explore how existing regulations impact the development of AI systems.

Explore our AI Playbook, a blueprint for developing and deploying AI, at www.mmventures.com/research.

AI tackles profound technical challenges

AI is significant because it successfully tackles a profound set of technical challenges. Increasingly, human capabilities – understanding, reasoning, planning, communication and perception – can be undertaken by software, at scale and at low cost. General analytical tasks, including finding patterns in data, that have been performed by software for many years can also be performed more effectively using AI.

Together, these capabilities create new opportunities in most business processes and consumer applications.

AI research is focused on five fields

Since its inception in the 1950s, AI research has focused on five fields of enquiry:

1. Knowledge: The ability to represent knowledge about the world. For software to possess knowledge, it must understand that certain entities, facts and situations exist in the world; these entities have properties (including relationships to one another), and these entities and properties can be categorised.

2. Reasoning: The ability to solve problems through logical reasoning. To reason is to apply logic to derive beliefs, related ideas and conclusions from information. Reasoning may be deductive (derive specific conclusions from general premises believed to be true), inductive (infer general conclusions from specific premises) or abductive (seek the simplest and most likely explanation for an observation).

3. Planning: The ability to set and achieve goals. For software to be able to plan, it must be capable of specifying a future, desirable state of the world and a sequence of actions enabling progress towards it.

4. Communication: The ability to understand written and spoken language. To communicate with people, software must have the ability to identify, understand and synthesise written or spoken human language.

5. Perception: The ability to make deductions about the world based on sensory input. To perceive, software must be able to organise, identify and interpret visual images, sounds and other sensory inputs.

Increasingly, human capabilities
– understanding, reasoning,
planning, communication and
perception – can be undertaken
by software, at scale and
low cost.

Progress in AI has unlocked new possibilities

Because most business processes and consumer applications involve knowledge management, reasoning, planning, communication or perception, progress in AI has unlocked significant new capabilities.

AI is enabling new possibilities				
Knowledge	Medical diagnosis	Drug creation	Media recommendation	
	Financial trading	Information synthesis	Consumer targeting	
Reasoning	Legal analysis	Asset management	Application processing	
	Games	Autonomous weapons	Compliance	
Planning	Logistics	Fleet management	Navigation	
	Network optimisation	Predictive maintenance	Demand forecasting	
Communication	Voice control	Intelligent agents	Customer support	
	Real-time transcription	Real-time translation	Client service	
Perception	Autonomous vehicles	Medical imaging	Authentication	
	Augmented reality	Surveillance	Industrial analysis	

Source: MMC, Veritas

In the following chapter, we describe specific AI use cases in eight sectors.

The applications of AI in industry are numerous and tangible

AI is not a solution seeking a problem; it is a tangible set of capabilities unlocking revenue growth and cost savings.

The capabilities of AI – its power to incorporate broader data sets into analyses, identify concepts and patterns in data better than rules-based systems, and enable human-to-machine conversation – have applications in all sectors and numerous business processes. In approximately 60% of occupations, at least 30% of constituent activities are technically automatable by adapting currently proven AI technologies (McKinsey Global Institute). As such, AI is a key ‘enabling technology’.

In approximately 60% of occupations, at least 30% of constituent activities are technically automatable by adapting currently proven AI technologies.

(McKinsey Global Institute)

Data-centric sectors will see the greatest impact

AI is being deployed in all sectors and to a wide variety of business processes. However, AI will have more numerous applications and greater impact in certain sectors.

AI’s impact will be greatest in sectors in which a large proportion of time is spent collecting or synthesising data, or undertaking predictable physical work. In several sectors (Fig. 17), professionals spend one third or more of their time on the above (McKinsey, Julius Baer).

These sectors include:

- Finance and insurance (50% of time)
- Retail, transport and trade (40% of time)
- Professional services (37% of time)
- Manufacturing (33% of time)
- Healthcare (33% of time)

Applications will be more limited in sectors in which data synthesis and processing activities are limited, or in which the majority of people’s time is spent managing others or undertaking unpredictable physical work. Occupations such as management and teaching will be more resistant to AI in the medium term.

Fig 17. In several sectors, professionals spend a third or more of their time collecting or synthesising data



Core use cases vary by sector

Use cases for AI are proliferating as understanding of the technology improves. Below, we describe 3) core

AI use cases in eight sectors: asset management, healthcare, insurance, law & compliance, manufacturing, retail, transport and utilities.

Sector	Core use cases:			
Asset Management	Investment strategy	Portfolio construction	Risk management	Client service
Healthcare	Diagnostics	Drug discovery	Monitoring	
Insurance	Risk assessment	Claims processing	Fraud detection	Customer service
Law & compliance	Case law	Discovery and due diligence	Litigation strategy	Compliance
Manufacturing	Predictive maintenance	Asset performance	Utility optimisation	
Retail	Customer segmentation	Content personalisation	Price optimisation	Churn prediction
Transport	Autonomous vehicles	Infrastructure optimisation	Fleet management	Control applications
Utilities	Supply management	Demand optimisation	Security	Customer experience

Source: M&I Ventures

Asset management

AI's ability to extract content from unstructured data using natural language processing, find subtle patterns in disparate data sets, and enable machine-to-human communication via chatbots, has multiple applications in asset management.

Core use cases include investment strategy, portfolio construction, risk management and client service.

By augmenting or automating many of an asset manager's tasks, AI enables asset managers to deliver to the mass affluent a degree of personalisation and service quality previously reserved for high net worth clients. Additionally, AI can decrease the cost of portfolio construction while improving quality – the era of the "robo-advisor".

Investment strategy: AI can improve a firm's investment strategy by synthesising its research and data, and incorporating broader data sets including unstructured information. Superior pattern recognition can then deliver better multi-objective optimisation. AI can balance a diverse range of inter-connected objectives (including fund deployment, risk and profitability) to enhance returns more effectively than rules-based systems.

Portfolio construction: AI tools can augment, and increasingly automate, an asset manager's process of portfolio construction. AI – "robo-advisors" – can analyse a client's goals,

and within a firm's investment rules develop personalised, optimised portfolios at low cost and high speed.

Risk management: AI can improve risk management by incorporating broader data sets and improving analytical processing. 90% of data generated today is unstructured information, stored outside traditional databases (International Data Group). Natural language processing enables additional data sets to be incorporated into firms' analyses. Other AI techniques, including deep learning, then enable patterns in data to be identified with greater granularity and confidence. Together, these capabilities enable risks to be identified and quantified more effectively.

Client service: Chatbot interfaces are being applied within and beyond asset management firms. Deployed in client-facing channels, natural language systems enable client enrolment, support and self-service. Embedded in internal tools, chatbots let account managers query client details and understand developments relevant to a client's portfolio in seconds instead of minutes. Fewer account managers can then provide a higher quality service to a greater number of clients.

Healthcare

In the next decade, AI can unlock a paradigm shift in healthcare to improve patient care and process efficiency. Automated diagnosis was an early use case for rudimentary AI in the 1980s. 'Expert systems' mimicked human approaches to diagnosis, applying rules-based inferences to bodies of knowledge. Modern AI, particularly deep learning, is more effective and applicable to a wider range of processes. Key use cases include diagnosis, drug discovery and patient monitoring.

Diagnosis: Deep learning systems can replace complex, human-coded sets of probabilistic rules and identify subtle correlations between vast, multi-variate data sets to deliver scalable, automated diagnosis. While systems are nascent, accuracy is improving rapidly. Separately, computer vision solutions powered by deep learning are transforming diagnostic imaging. While human radiologists require extensive expertise and years of training to identify abnormalities in magnetic resonance images and ultrasounds, deep learning systems trained on large data sets deliver impressive results. Diagnostic imaging, powered by deep learning, now offers human-level accuracy and high speed in select contexts.

Drug discovery: Today's drug discovery process is lengthy, averaging 12 years to market (California Biomedical Research Association). Expense and uncertainty are also prohibitive; drug development costs an average of \$359m and just 2% of US preclinical drugs are approved for human use (California Biomedical Research Association). AI is being applied to multiple stages of the drug development process to accelerate time to market and reduce uncertainty. AI is being applied to synthesise information and offer hypotheses from the 10,000 research papers published daily, predict how compounds will behave from an earlier stage of the testing process, and identify patients for clinical trials.

Monitoring: Monitoring the vital signs of patients on non-acute wards, or at-risk individuals in the home, remains a manual process undertaken periodically. AI can synthesise signals from inexpensive wearable devices worn by patients to deliver clinical-grade monitoring, and enable a large group of patients to be monitored in real-time by a single nurse. As data sets are amalgamated and algorithms are tuned, AI will offer predictive analytics. Patients on a ward, or at home, who require further hospital care can be identified and supported, and unnecessary use of hospital beds reduced.

\$359m
The average cost of drug development: \$359m. Just 2% of US pre-clinical drugs are approved for human use.

Source: California Biomedical Research Association

Insurance

While the fundamentals of insurance – customer prospecting, risk assessment, claims processing and fraud detection – have remained unchanged, modern AI can improve every stage in the insurance process to deliver efficiency savings and improved customer experience. By identifying patterns in data better than rules-based systems, AI can improve and accelerate decision-making and claims processing, reduce fraud and automate a large proportion of customer service enquiries.

Risk assessment: AI can gather information from broader data sets, including web and social media profiles, to compile richer customer information and inform risk assessment. AI can then assess the risk of individual policies more accurately than rules-based systems, by detecting non-linear patterns in multi-variate data sets and making more accurate projections.

Claims processing: AI can reduce time-to-quote, time-to-claim and claims processing costs for consumers and insurers. To accelerate claims processing, AI systems can automatically extract and classify structured and unstructured information from insurance policies and claim forms. By analysing images of damaged assets, computer vision systems can automatically classify claims. Through improved pattern recognition applied to prior cases, AI can also predict settlement costs. Algorithms using deep learning are effective for image analysis, while Bayesian (probability-based) AI is useful for predicting settlement costs.

Fraud detection: Insurance fraud costs UK insurers £1.3bn annually and adds £50 to the average policyholder's annual bill (Association of British Insurers). UK insurers invest over £200m annually to tackle the challenge (Association of British Insurers). Fraud detection algorithms enhanced with AI can identify fraudulent transactions, while reducing false positives, more effectively than traditional approaches.

Customer service: Chatbot interfaces integrated with insurers' databases can use natural language processing to offer 24/7 product information and answers to policyholders' enquiries, in a scalable, inexpensive and personalised channel.

Fraud detection algorithms enhanced with AI can identify fraudulent transactions, while reducing false positives, more effectively than traditional approaches.

Law and Compliance

AI's abilities to process language in documents, synthesise knowledge and automate reasoning have broad application in the legal services and compliance sector. With junior lawyers spending a high proportion of their time accessing and collating information, scope for augmentation and automation is considerable. Key AI use cases include identifying relevant

case law, processing documents for discovery and due diligence, and informing litigation strategy. In October 2018, Harvard Law School Library advanced its 'Caselaw Access Project' by releasing over 40 million pages of digitised legal information, including every reported state and federal legal case in the United States from the 1600s to the summer of 2017

- providing extensive further data to train AI systems.

Regarding compliance, costs have grown significantly since 2008 – particularly for financial services firms. With rules-based software poorly suited to catching infractions, banks have invested in additional compliance personnel. Citl, while reducing its global headcount 32% between 2008 and 2016, doubled its regulatory and compliance staff to 29,000 – over 13% of its workforce (Citl). AI's ability to learn patterns of behaviour over time, and highlight unusual activity in real-time, offers greater scalability at lower cost.

Case law, discovery and due diligence: Natural language processing AI can identify, classify and utilise content from databases and unstructured documents at scale and speed, saving legal firms time and cost for routine document review. Use cases include sourcing and ranking relevant case law and identifying key documents in due diligence and discovery processes. With a merger and acquisition data room containing an average of 34,000 pages for review (Luminance), AI can increase business velocity and reduce costs.

Litigation strategy: AI can analyse past judgements at greater speed, granularity and subtlety than has been possible previously. By anticipating the probability of different outcomes, lawyers can inform and enhance their strategic decision-making. In high volume areas, such as personal injury, software can help a firm decide whether to accept a case. In high value areas, including corporate litigation, software can suggest the probability of a particular outcome based on judges' prior behaviour and opposing lawyers' tendency to settle or proceed to trial.

Compliance: Preventing accidental or deliberate breaches of policy, from the theft of sensitive data to accidentally misaddressing an email containing a customer database, is challenging for rules-based systems. By learning the habits of users over time, AI systems can flag potential compliance breaches in real-time, before they occur, with sufficient accuracy to enable broad deployment.

AI's ability to process language in documents, synthesise knowledge and automate reasoning has broad application in the legal services and compliance sectors.

Manufacturing

AI can significantly improve manufacturers' efficiency and profitability. Overall Equipment Effectiveness (OEE), a measure of manufacturers' productivity relative to potential, varies widely by industry, from 75%-91% (IHS Research). The performance of companies within the same industry also varies widely, offering scope for competitive advantage. AI can boost OEE and profitability by predicting equipment failure (to reduce unplanned downtime), improving assets' operational efficiency, and reducing utility supply costs.

Predictive maintenance: The failure of production assets is costly; one hour of unplanned downtime on an automotive assembly line can cost a manufacturer \$1.5m (MMC Ventures). AI can identify subtle patterns in data from vibration, temperature, pressure and other sensors to identify leading indicators of equipment failure. By predicting more accurately which components are likely to fail, and when, parts can be proactively replaced to prevent failures and save money.

Asset performance: AI can improve the operation of high value assets, including gas and wind turbines, to optimise yield. Rules-based programs deliver limited results when applied to complex tasks, such as tuning fuel valves on a gas turbine to optimise combustion while reducing wear and emissions. Applying neural networks to optimise the turbine inputs can improve results by 20% or more.

Utility optimisation: Optimising the purchase and consumption of utilities, such as power and water, according to real-time demands on a factory floor is too challenging and variable to manage using rules-based software. AI enables companies to anticipate, and align, utility consumption with process requirements in real-time, lowering utility consumption by 5% or more.

Retail

E-commerce, now 17% of UK retail sales and growing (eMarketer), has transformed the quantity, breadth and granularity of data available to retailers. Retailers that turn data into insight can increase competitive advantage by engaging, monetising and relating customers more effectively. Every stage of a retailer's customer journey – from lead generation and content selection to price optimisation and churn prediction – can be improved by AI algorithms that ingest richer data sets and identify patterns in them better than rules-based systems. By enabling analytics at the 'per-customer' level, AI is introducing the era of retail personalisation. Leaders enjoy competitive advantage; 75% of Netflix users select films recommended to them by the Company's AI algorithm (Netflix).

Customer segmentation: Limitations in available data, and the linear analysis of information, inhibit the ability of traditional customer segmentation software to identify desirable customer attributes. Deep learning algorithms enable natural language processing, which enables retailers to access additional data sets including social media data. Deep learning algorithms also offer more granular analysis than rules-based systems, to optimise segmentation, channel selection and messaging.

Content personalisation: Most content presented to online shoppers is irrelevant or poorly suited to users' preferences, reducing conversion to an average of 1.10% on smartphones and 2.8% on desktops (Adobe). As with customer segmentation, AI offers additional unstructured data sets for analysis, and improved multivariate analysis to identify more subtle correlations than rules-based systems can detect. When Netflix recommends content to a user, in addition to analysing a user's actions, ratings and searches, the Company's AI algorithm considers social media data and meta-data from third parties. The Company is now analysing images from content, including colour palette and scenery, for deeper personalisation.

Price optimisation: A 1% change in price provides, on average, a 10% change in profitability (BlueYonder). The smaller a company's margins, the greater the impact. Willingness to pay is a key determinant for price. AI enables price optimisation (that is more sophisticated than traditional 'cost plus', 'relative-to-competitors' or 'odd pricing' (EO 99) models. By identifying correlations within and between data sets, AI can better

75%

of Netflix users select films recommended to them by the Company's AI algorithms.

Source: Netflix

optimise for factors including price elasticity, revenue, profit, product availability and phases in a product's lifecycle (introduction or end-of-life).

Churn prediction: Traditional programs struggle to incorporate new sources of information, maximise the value from multi-variate data sets or offer granular recommendations. AI-powered churn prediction can identify leading indicators of churn more effectively, and improve remediation by predicting more accurately the format and content of successful interventions.

Transport

The transport sector will be transformed by AI. Breakthroughs in computer vision are enabling the age of autonomous vehicles – self-driving cars, buses and trucks. The implications, from shifts in sector value chains to new business models, will be profound (see Chapter 8). In addition to enabling autonomy, AI can be applied to the many prediction and optimisation challenges – from congestion modelling to fleet management – at the core of today's logistics networks.

In addition to enabling autonomy, AI can be applied to the many prediction and optimisation challenges – from congestion modelling to fleet management – at the core of today's logistics networks.

Autonomous vehicles: AI computer vision systems enable vehicles to sense and identify the physical features and dynamics of their environment, from road lanes to pedestrians and traffic lights, with a high degree of accuracy. Combined with AI data processing and planning algorithms, AI is enabling the age of autonomous transport. Cars, buses and trucks will be able to operate and guide themselves, without human involvement. SAE International, a US-based global professional association and standards body, has identified five degrees of vehicle autonomy, from Level 0 (no automation) to Level 5 (full automation; no requirement for human control). Select companies, including Google, intend to release vehicles offering Level 5 automation. Challenged by the autonomous vehicle programmes of Google, Uber and Tesla, incumbent

2021
The date by which private hire taxi firm Addison Lee intends to deploy self-driving cars in London.

Source: Addison Lee

manufacturers are accelerating their own initiatives by increasing investment and making acquisitions. Ford intends to deliver high-volume availability of at least a Level 4 autonomous vehicle by 2021. In October 2018, in the UK, private hire firm Addison Lee announced its intention to deploy self-driving cars in London by 2021, by partnering with UK autonomy company Oxbotica.

Infrastructure and system optimisation: AI's abilities to detect patterns and optimise complex data are being applied to traffic, congestion and infrastructure challenges in transport systems. Predicting traffic flows, or modelling the deterioration of transport infrastructure, are difficult because inputs are complex (combining traffic, construction and environmental data) and because the relationships between inputs and outputs are non-linear (Transportation Research Circular). In these contexts, machine learning and deep learning systems are well suited to deliver better results than rules-based systems.

Fleet management: Transportation fleets are pervasive, from the logistics networks that underpin the economy to taxi fleets and food delivery services that provide point-to-point convenience. AI can optimise pick-ups, route planning and delivery scheduling to maximise asset utilisation, while considering economic, social and environmental impacts.

Control applications: Machine learning systems are well suited to the numerous prediction and optimisation challenges presented by air traffic control, vehicle traffic signalling, and train control.

Utilities

Information processing will become critical to utility companies, and their business models, as the utility sector undergoes a greater change in the next 25 years than it has during the previous 150. 'Prosumers' – consumers who also own capacity for energy production – will require integration into the energy market. By processing data more intelligently, AI will be a significant value driver in this transition. AI use cases for utility companies are varied, from demand optimisation and security to customer experience.

The foundations for AI adoption in the utilities sector are robust. 67% of utility companies – a higher proportion than in any other sector – use 'internet of things' (IoT) technologies such as sensors (Gartner). Further, compared with peers in other sectors, utility CIOs have a stronger focus on cost reduction, managing geographically dispersed assets and security.

Supply management: AI algorithms can predict changes in supply, including those caused by the intermittency of renewable resources, more effectively than rules-based systems – enabling smaller reserves and greater cost savings. AI solutions can also optimise supply networks, which are becoming increasingly complex as consumers deploy sources of renewable energy that contribute energy back to the National Grid.

Demand optimisation: By identifying detailed patterns in consumer behaviour, AI algorithms can move consumption of energy from periods of peak use and high prices to times of lower demand and cost.

Security: Rules-based systems struggle to deliver system security given the continually evolving nature of security threats. By identifying abnormal patterns in network behaviour, deep learning systems can identify breaches in network security that elude traditional programs.


Customer experience: Chatbots, which offer natural language conversations powered by deep learning algorithms, offer consumers self-service account administration, product information and customer service.

In Chapter 8, we explore the profound implications of the proliferation of AI across multiple sectors.

67%

of utility companies – a higher proportion than in any other sector – use 'internet of things' (IoT) technologies such as sensors.

Source: Gartner



Chapter 3




Why has AI come of age?

Specialised hardware, availability of training data, new algorithms and increased investment, among other factors, have enabled an inflection point in AI capability. After seven false dawns since the 1950s, AI technology has come of age.

Summary

- After seven false dawns since its inception in 1956, AI technology has come of age.
- The capabilities of AI systems have reached a tipping point due to the confluence of seven factors: new algorithms; the availability of training data; specialised hardware; cloud AI services; open source software resources; greater investment; and increased interest.
- Together, these developments have transformed results while slashing the difficulty, time and cost of developing and deploying AI.
- A virtuous cycle has developed. Progress in AI is attracting investment, entrepreneurship and interest. These, in turn, are accelerating progress.



Recommendations

Executives

- Recognise that AI technology has come of age and will be a key enabler, and potential threat, in the coming decade.
- Pioneers are deploying AI at an accelerating rate. Familiarise yourself with the dynamics of enterprise AI adoption (Chapter 4).
- Explore the many applications of AI (Chapter 2), and AI's implications (Chapter 8), to lead and contribute to AI initiatives in your organisation.

Entrepreneurs

- AI technology can deliver tangible benefits today. Seek opportunities to incorporate AI within your software, where appropriate, whether or not you are an 'AI company'.
- Familiarise yourself with the latest developments in AI technology (Chapter 5) and talent (Chapter 6) to enable your AI initiatives.

Investors

- AI will be a powerful enabler for portfolio companies – and a threat. Evaluate whether portfolio companies are embracing AI as a means of competitive advantage.
- With AI technology at a tipping point, seek opportunities to invest directly or indirectly in companies taking advantage of AI.
- Explore recent developments in AI technology (Chapter 5) to identify emerging areas of opportunity.

Policy-makers

- Review policymakers' key initiatives and identify opportunities for further sector support. In the UK, key programmes and studies include: the UK Government's £1bn 'AI sector deal'; recommendations from the House of Lords Select Committee on AI's (VAI in the UK: ready, willing and able?); and findings of the All-Party Parliamentary Group on AI.

Explore our AI Playbook, a blueprint for developing and deploying AI, at www.nimcventures.com/research.

There are seven enablers of AI

Research into AI began in 1956. After seven false dawns, in which results from unsophisticated systems fell short of expectations, AI capability has reached a tipping point.

AI is now delivering significant utility and its abilities are advancing rapidly.

AI capabilities have been transformed in the last four years due to:

1. the development of improved AI algorithms;
2. increased availability of data to train AI systems;
3. specialised hardware to accelerate training of AI algorithms;
4. cloud-based AI services to catalyse developer adoption;
5. open source AI software frameworks that enable experimentation;
6. increased investment in AI by large technology companies and venture capitalists;
7. greater awareness of AI among investors, executives, entrepreneurs and the public.

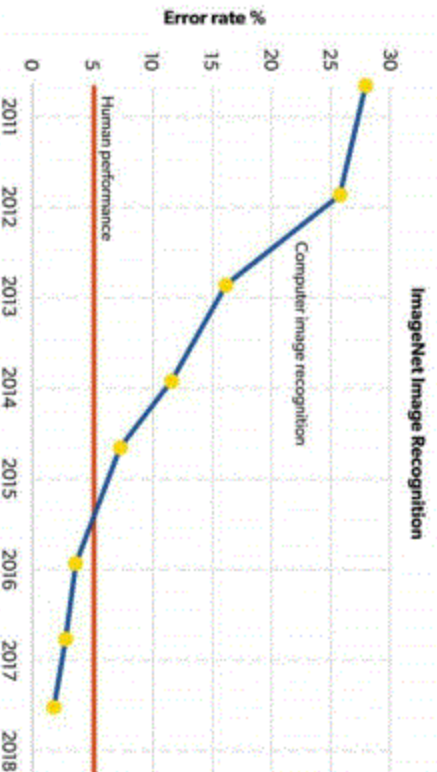
Together, these developments have improved results from AI systems and increased the breadth of challenges to which they can be applied. They have also irreversibly reduced the difficulty, time and cost of developing basic AI systems.

1. Enhanced algorithms provided improved results

Deep learning, a fruitful form of machine learning, is not new; the first specification for an effective, multilayer neural network was published in 1965. In the last decade, however, evolutions in the design of deep learning algorithms have transformed results, delivering breakthrough applications in areas including computer vision (Fig. 18) and language (Fig. 19).

Convolutional Neural Networks (CNNs) have dramatically improved computers' ability to recognise objects in images. Employing a design inspired by the visual cortex of animals, each layer in a CNN acts as a filter for the presence of a specific pattern. In 2015, Microsoft's CNN-based computer vision system identified objects in pictures more effectively (95.1% accuracy) than humans (94.9% accuracy) (Microsoft). In the last 36 months, performance has improved further (Fig. 18). Broader applications of CNNs include video classification and speech recognition.

Fig 18. Convolutional neural networks are delivering human-level image recognition




Source: <https://www.aifeng.ai/>

Recurrent Neural Networks (RNNs) are delivering improved results in speech recognition and beyond. While data progresses in a single direction in conventional ("feed forward") neural networks, RNNs have feedback connections that enable data to flow in a loop. With additional connections and memory cells, RNNs "remember" data processed thousands of steps ago and use it to inform their analysis of what follows. This is valuable for speech recognition, where interpretation of an additional word is enhanced by analysis of preceding ones.

Microsoft reported that its speech recognition system achieved human-level recognition for the first time in history. Improvements are continuing.

The Long Short-Term Memory (LSTM) model is a particularly effective recent RNN architecture. From 2012, Google used LSTMs to power speech recognition in the Android platform. In October 2016, Microsoft reported that its LSTM speech recognition system achieved a word error rate of 5.9% – human-level speech recognition for the first time in history (Microsoft) (Fig. 19). By August 2017, word error rate had been reduced to 5.1% (Microsoft). Improvements are continuing.

Fig 19. Recurrent neural networks are delivering human-level speech recognition



Source: <https://www.ft.com/g8>

2. Extensive data enabled AI systems to be trained

Training neural networks typically requires large volumes of data – thousands or millions of examples, depending on the domain. The creation and availability of data has grown exponentially in recent years, enabling AI.


Today, humanity produces 2.5 exabytes (2,500 million gigabyte) of data daily (Google). 90% of all data has been created in the last 24 months (SINTEF). Data has ballooned as humanity passed through two waves of data creation, and now enters a third.

The first wave of data, beginning in the 1980s, involved the creation of documents and transactional data. It was catalysed in the 1990s by the proliferation of internet-connected desktop PCs. Then, in the 2000s and 2010s, pervasive, connected smartphones drove a second wave of data with an explosion of unstructured media (email, photos, music and video), web data and metadata.

Today we enter the third age of data. Machine sensors deployed in industry and the home provide additional monitoring – analytical- and meta-data. With much data created today transmitted for use via the internet, growing internet traffic is a proxy for humanity's increasing data production. In 1992, humanity transferred 100GB of data per day. By 2020, we will transfer 61,000GB per second (Fig. 20) (Cisco, MMC Ventures).

The development of AI has been catalysed further by the creation of specialist data resources. ImageNet, a free database of 14.2 million hand-labelled images, has supported the rapid development of deep learning algorithms used to classify objects in images.

Fig 20. Global Internet traffic is increasing exponentially, reflecting growth in data production



Source: Cisco, MMC Ventures

3. Specialised hardware accelerated AI system training

Graphical Processing Units (GPUs) are specialised electronic circuits that slash the time required to train the neural networks used in deep learning-based AI.

Modern GPUs were developed in the 1990s, to accelerate 3D gaming and 3D development applications. Panning or zooming a camera in a simulated 3D environment uses a mathematical process called a matrix computation.

Microprocessors with serial architectures, including the Central Processing Units (CPUs) that interpret and execute commands in today's computers, are poorly suited to the task.


GPUs were developed with massively parallel architectures (NVIDIA's GeForce RTX 2080 TI GPU has 4,352 cores) to perform matrix calculations efficiently. Training a neural network involves numerous matrix computations. GPUs,

while conceived for 3D gaming, therefore proved ideal for accelerating deep learning.

A simple GPU can increase five-fold the speed at which a neural network can be trained. Ten-fold or larger gains are possible. When combined with Software Development Kits (SDKs) tuned for popular deep learning frameworks, even greater improvements can be realised. In a 36 month period beginning in 2013, successive GPUs and SDKs enabled a 50x increase in the speed at which certain neural networks could be trained (Fig. 21).

In the last 36 months, advances in AI technology are creating new possibilities. Custom silicon, designed from inception for AI, is enabling a new generation of AI accelerators (Chapter 5).

Fig 21. GPUs enabled neural networks to be trained 50x faster



Adapted from a report published on 20 January. Source: NVIDIA

4. Cloud AI services fuelled adoption

Leading cloud technology providers including Google, Amazon, IBM and Microsoft offer cloud-based AI infrastructure and services, catalysing developers' use of AI.

The providers' infrastructure platforms include environments in which to develop and deploy AI algorithms, and GPUs-as-a-service, to power them.

Their services comprise a burgeoning range of on-demand AI capabilities, from image recognition to language translation, which developers can incorporate in their own applications. Google Machine Learning offers application programming interfaces (APIs) for: computer vision (object identification, explicit content detection, face recognition and image sentiment analysis); speech (speech recognition and speech-to-text); text analysis (entity recognition, sentiment analysis, language detection and translation) and more. Microsoft Cognitive Services include over 24 services in the fields of vision, speech, language, knowledge and search.

The accessibility and relative affordability of cloud providers' AI infrastructure and services are significantly increasing adoption of AI among developers.


5. Open source software catalysed experimentation

The release of open source AI software frameworks has lowered barriers to entry for experimentation and proficiency in AI.

Researchers, and providers of cloud infrastructure and AI services who benefit from the proliferation of AI and data-intensive applications, have open-sourced AI frameworks and libraries of algorithms to catalyse developers' adoption of AI. Popular open source platforms include TensorFlow (Google), Caffe2 (Facebook), Cognitive Toolkit (Microsoft), TorchNet (Facebook), H2O (H2O.ai) and Mahout (Apache Software Foundation).

Each framework offers benefits. Caffe2 is a scalable deep learning framework that processes images at speed. Cognitive Toolkit provides high performance on varying hardware configurations. H2O reduces time-to-value for AI-powered enterprise data analysis. Mahout provides scalability and pre-made algorithms for tools such as H2O. Google's decision to open source TensorFlow in November 2015 was particularly significant, given the software's sophistication. Engagement with TensorFlow has been rapid (Fig. 22). Within two years, the framework has attracted 30,000 developer commitments and 80,000 stars on GitHub, where developers store projects (Google).

Fig 22. Engagement with the TensorFlow framework has been significant and rapid



5. Investment in AI increased fifteen-fold

Venture capital firms are investing aggressively in AI, given scope for value creation. Investment dollars into early stage AI companies globally have increased fifteen-fold in five years (Fig. 23), to an estimated \$15bn in 2018 (CB Insights, MMC Ventures).

Today's leading technology companies – including Apple,

Amazon, Facebook, Google, IBM, Microsoft and Salesforce

– are also spending heavily on research and personnel to

develop and deploy AI. Internal corporate investment on AI,

among just the top 35 high tech and advanced manufacturing

companies investing in AI, may be 2.0x to 4.5x greater than the


capital invested by venture capital firms, private equity firms

and other sources of external funding combined (McKinsey),

further catalysing progress.

Investment dollars into early stage AI companies globally have increased fifteen-fold in five years, to an estimated \$15bn in 2018.
(CB Insights, MMC Ventures)

Fig 23. Venture capital investment in AI has increased 15-fold in five years



Source: CB Insights, MMC Ventures

7. Awareness of AI has grown significantly

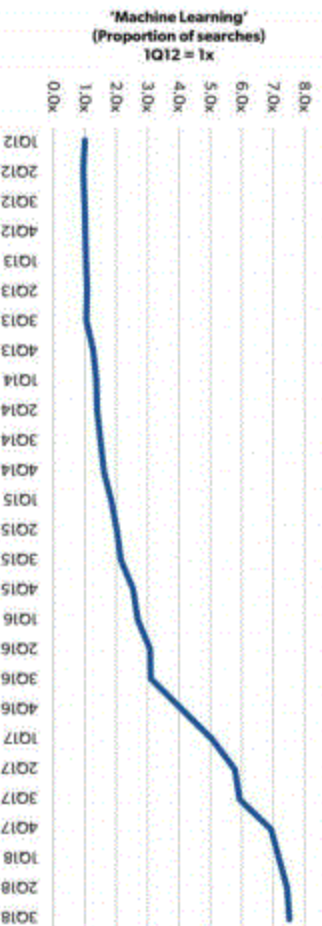
Public interest in AI, measured by the proportion of Google searches for 'machine learning', has increased more than seven-fold in six years (Fig. 24).

Executives' awareness of AI has grown following extensive coverage in business publications. In the last 12 months, 5,700 articles referencing AI have appeared in financial publications such as the Financial Times, Fortune, Investors Chronicle, The Wall Street Journal and Thomson Reuters, Bloomberg Businessweek, the Financial Times, Forbes, Fortune, the Harvard Business Review and The Wall Street Journal (Signal). One third of these references have appeared in the last 12 weeks.

In the popular press, whether relevant (the opportunities and threats posed by automation) or less so ('killer robots'), 21,800 articles in US and UK newspapers have referred to AI, fueling public interest (Signal).

Public interest in AI, measured by the proportion of Google searches for 'machine learning', has increased more than seven-fold in six years.

Fig 24. Interest in AI has increased 7-fold



Source: Google Trends

Chapter 4



The race for adoption

AI may be the fastest paradigm shift in technology history. Increasing adoption masks a growing divergence, among nations and within industries, between leaders and laggards.

Summary

- AI adoption has tripled in 12 months. One in seven large companies has adopted AI; in 24 months, two thirds of large companies will have live AI initiatives. In 2019, AI 'crosses the chasm' from early adopters to the early majority.

- AI may be the fastest paradigm shift in technology history. In the course of three years, the proportion of enterprises with AI initiatives will have grown from one in 25 to one in three. Adoption has been enabled by the prior paradigm shift to cloud computing, the availability of plug-and-play AI services from global technology vendors and a thriving ecosystem of AI-led software suppliers.

- Great expectations are fueling adoption. Executives expect AI to have a greater impact than any other emerging technology, including Blockchain and IoT.

- Increasing overall adoption masks a growing divergence between leaders and laggards. Leaders are extending their advantage by learning faster and increasing investment in AI at a greater pace than laggards.

- Globally, China leads the race for AI adoption. Twice as many enterprises in Asia have adopted AI, compared with companies in North America, due to government engagement, a data advantage and fewer legacy assets.

- Sector adoption is uneven and in a state of flux. 'Early adopters' (financial services and high-tech companies) maintain a lead while 'movers' (retail, healthcare and media) are rapidly catching up. Government agencies, education companies and charities are laggards in AI adoption. Vulnerable members of society may be among the last to benefit from AI.

- AI is advancing across a broad front. Enterprises are using multiple types of AI application, with one in ten enterprises using ten or more. The most popular use cases are chatbots, process automation solutions and fraud analytics. Natural language and computer vision AI underpin many prevalent applications as companies embrace the ability to replicate traditionally human activities in software for the first time.

- Leaders and laggards face different adoption challenges. Laggards are struggling to gain leadership support for AI and to define use cases. Leaders' difficulties, in contrast, have shifted from 'if' to 'how'. Leaders are seeking to overcome the difficulty of hiring talent and address cultural resistance to AI.

- AI initiation has shifted from the C-suite to the IT department. Two years ago, CIOs initiated two thirds of AI initiatives. In 2019, as corporate engagement with AI shifts from 'if' to 'how', the IT department is the primary driver of projects.

- Companies prefer to buy, not build. AI. Nearly half of companies favour buying AI solutions from third parties, while a third intend to build custom solutions. Just one in ten companies are prepared to wait for AI to be incorporated into their favourite software products.

- Workers expect AI to increase the safety, quality and speed of their work. As companies' AI agendas shift from revenue growth to cost reduction initiatives, however, workers are concerned about job security.

Recommendations

Executives

- With AI 'crossing the chasm' to the early majority, the time to act is now. Develop an AI strategy to avoid losing competitive advantage.
- AI leaders are extending their advantage with increasing investment in AI. Ensure AI initiatives are a budget priority to enable test-and-learn deployments.
- Identify and address barriers to adopting AI within your organisation. Are they challenges of 'if' (lack of leadership support, difficulty defining use cases) or 'how' (attracting AI talent, cultural concerns)?
- Initiation of AI projects has shifted from the C-suite to IT departments and lines of business. Support these teams' efforts to catalyse AI.
- If buying AI, explore the ecosystem of 1,600 early stage software companies in Europe that have AI at the heart of their value proposition (Chapter 7).
- Staff may be concerned about job security. Engage with employees to explain how AI can augment their roles.

Entrepreneurs

- Buyers are diverging into leaders and laggards, and 'buyers' versus 'builders'. Qualify attractive prospects early in your engagement process and align the benefits you describe with pain points typical for the buyer persona.
- A quarter of buyers wish to buy an AI solution before customising it further to their industry requirements. Be prepared to iterate your offering in accordance with buyers' needs, in return for access to data and public endorsement.
- When developing go-to-market plans and messaging, be mindful of significant differences in departments' interest in AI.
- With initiation of AI projects shifting from C-suites to IT teams and department heads, optimise your engagement plans to mitigate these groups' concerns.

Investors

- Shifts in sector adoption present new areas of opportunity and change the go-to-market dynamics for startups in existing segments. Consider the implications of sectors and departments at a tipping point in AI adoption.
- Growing adoption of AI presents a new backdrop for the efforts of early stage AI companies. Assess whether prospects and portfolio companies are developing the competencies required to sell to a maturing market.

Policy-makers

- 17 countries have national AI strategies. To avoid falling behind, countries must challenge the ambition and scope of their strategies, while supporting their implementation with increased investment, expanded plans for the cultivation of talent and extended strategies for access to data.
- The government sector, and non-profit organisations, are laggards in AI adoption. Redouble efforts to increase public sector and non-profit organisations' use of AI, given the benefits AI can deliver.

AI adoption has tripled in 12 months

Large companies are adopting AI at a rapidly accelerating rate. Just 4% of enterprises had adopted AI 12 months ago (Gartner). Today, 14% of enterprises have deployed AI. A further 23% intend to deploy AI within the next 12 months. Adoption will continue to accelerate: in two years, nearly two thirds of large companies will have live AI initiatives (Fig. 25).


AI deployment is proliferating as:

- Widespread awareness of AI drives a growing volume of enterprise test-and-learn initiatives;
- Early proof-of-concept projects mature, demonstrating value and catalysing further investment;
- Understanding of AI, although low, is improving and driving investment;
- Maturing AI technology – and a burgeoning range of inexpensive or open source AI APIs, frameworks and tooling – lower barriers to entry. Enterprises can achieve more with AI faster, cheaper and with less expertise than 24 months ago;

- Enterprises mitigate skills shortages by recruiting chief science officers, researchers, data scientists and machine learning engineers – and up-skilling existing employees;
- Enterprises embrace a rich ecosystem of ‘best-of-breed’ third-party AI software suppliers. Europe is home to over 1,600 innovative, early stage software companies with AI at the heart of their value proposition (Chapter 7). Serving a broad range of sectors and business functions, they provide an accessible ‘on-ramp’ to AI with superior results and rapid time-to-value.

Today, 14% of enterprises have deployed AI. A further 23% intend to deploy AI within the next 12 months.
(Gartner)

Fig 25. One in seven large companies has deployed AI



Base: AI investing, n=2,832

Where are your organisation's plans in terms of artificial intelligence?

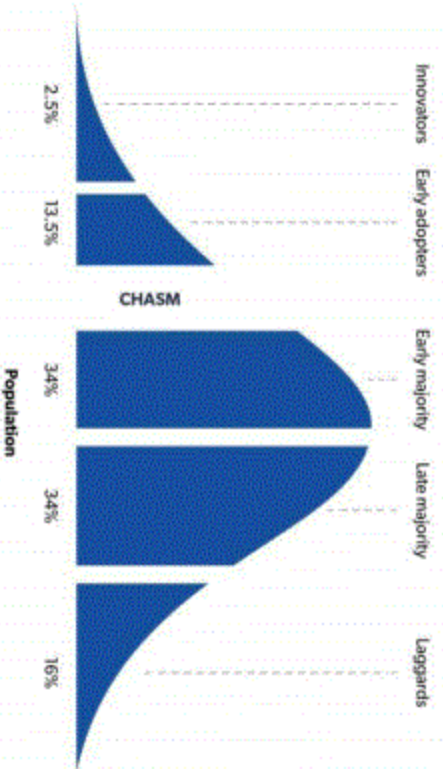
Source: Gartner, 2019 CIO Survey. CIOs Have Moved to the Importance of AI, Figure 1, 23 January 2019

In 2019, AI 'crosses the chasm' to the early majority

By the end of 2019, over a third of enterprises will have deployed AI. Adoption of AI has progressed extremely rapidly from innovators and early adopters to the early majority. By the end of 2019 AI will have 'crossed the chasm', from visionaries to pragmatists, at exceptional pace – with profound implications for companies, consumers and society.

Over three years, the proportion of companies with AI initiatives will have grown from one in 25 to one in three. (Gartner)

Fig 26. In 2019, AI 'crosses the chasm' to the early majority



Source: Forrester Research, Gartner Research

AI may be the fastest paradigm shift in technology history

AI may be the fastest major paradigm shift in the history of enterprise technology. In the course of three years, the proportion of companies with AI initiatives will have grown from one in 25 to one in three (Gartner).

Companies can enjoy initial benefits from AI with relative ease. Following the cloud computing revolution, and the emergence of a rich ecosystem, AI application providers (Chapter 6), enterprises can engage with 'best of breed' AI applications via the cloud to derive value from their data. They may also take advantage of 'plug and play' cloud AI services from global technology vendors including Amazon, Google, IBM and Microsoft.

While a deeper, structural embrace of AI – that may include hiring data scientists and re-mapping data pipelines – will require greater time and investment, the above factors are enabling the adoption of a new technology paradigm at unprecedented speed.

Great expectations are fuelling adoption

Adoption is being catalysed by companies' growing conviction in AI's potential. A greater proportion of executives believe AI will be a 'game changer' than any other emerging technology – including cloud, mobile, IoT, blockchain or APIs (Fig. 27).

Fig. 27 AI tops the list of technologies companies perceive as 'game-changing'

2019 CIO Agenda Which technology area do you expect will be a game-changer for your organisation?			
	Top performers (n = 230)	Typical performers (n = 2,329)	Trailing performers (n = 276)
Artificial Intelligence/Machine Learning	40%	25%	24%
Data Analytics (including Predictive Analytics)	23%	25%	21%
Cloud (including XaaS)	12%	10%	14%
Digital Transformation	10%	9%	7%
Mobile (including 5G)	7%	6%	5%
Robotic Process Automation (RPA)	6%	2%	1%
Internet of Things	6%	10%	11%
Blockchain	5%	4%	5%
Automation	3%	5%	5%
Information Technology	3%	2%	1%
APIs	2%	1%	0%
Immersive Experience	2%	1%	2%
Business Intelligence	2%	3%	5%
Cybersecurity	2%	1%	1%
Industry-Specific	2%	4%	5%
CRM	1%	2%	3%
ERP	1%	3%	3%

Source: Gartner (January 2019)

China leads the race in AI adoption

While adoption of AI has increased in all regions, companies in Asia/Pacific have been the most proactive in embracing AI. Twice as many enterprises in Asia/Pacific – one in five – have adopted AI today, compared with one in ten companies in North America (Gartner) (Fig. 28). Within Asia/Pacific, Chinese companies lead in AI adoption. Beijing, Shanghai, Guangdong, Zhejiang and Jiangsu are primary hubs. Further, the proportion of companies in Asia/Pacific with no interest in deploying AI – one in 14 – is half that of North America (Fig. 30).

“China’s rapid rise in AI has been a wake-up call for nations, industries and corporate executives globally.”

(MIT Sloan Management Review)

Chinese companies’ adoption of AI is being catalysed by:

1. **Government policy:** In 2017, China published its “Next Generation Artificial Intelligence Development Plan”.

A three-step plan for leadership in AI by China and Chinese companies, the roadmap seeks to establish Chinese competitiveness in AI by 2020, deliver breakthroughs in AI by 2025, and establish global leadership in AI by 2030.

2. **A data advantage:** AI systems typically improve by ingesting training data. Chinese companies have a dual advantage: more permissive policies than Europe regarding use of personal data, and less siloed data within companies. 78% of leading Chinese companies maintain their corporate data in a centralised data lake, compared with 37% of European and 43% of US pioneers (MIT Sloan Management Review).

3. **Fewer legacy assets:** Chinese companies typically have fewer legacy applications and processes, presenting opportunities to leapfrog European and American companies that have extensive existing systems and associated integration requirements.

Talent and personnel-related concerns are Chinese companies’ primary impediments to AI adoption. The AI talent pool in the United States is currently over 50% larger than in China (South China Morning Post). A greater proportion of pioneering Chinese companies – six in ten – highlight AI talent shortages than American or European enterprises (MIT Sloan Management Review). The impact of automation upon society is also a pressing concern for Chinese companies. Chinese companies have a greater focus on efficiency projects than revenue generating initiatives. As a result, two thirds of pioneering AI companies in China expect AI to reduce the size of their workforces, compared with a third of European peers.

Fig 28. “Deployed AI” (% of companies) – twice as many enterprises in Asia/Pacific than in North America have deployed AI



Source: Gartner

Chinese companies have a dual advantage: more permissive policies than Europe regarding use of personal data, and less siloed data within companies.

Fig 29. 'Deploying AI in the next 12 months'
(% of companies)



Fig 30. 'No interest in deploying AI' – the proportion
of companies in Asia/Pacific with no interest in
deploying AI (% of companies)

Source: Gartner

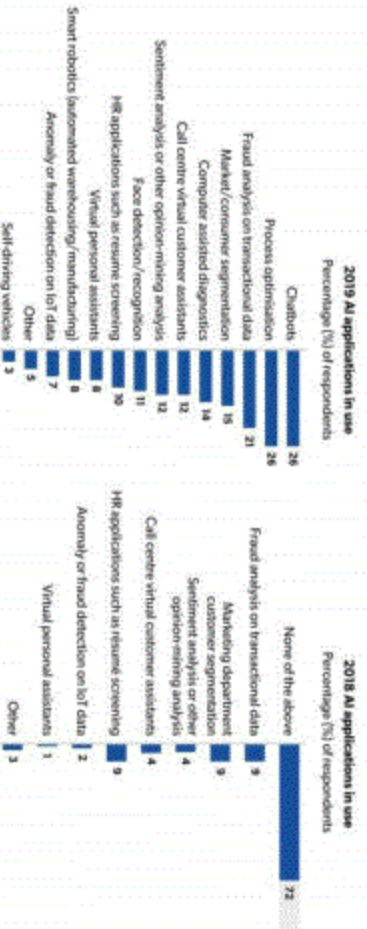
Source: Gartner

Use of AI applications is advancing across a broad front

Adoption is advancing not only substantially but across a broad front. (Fig. 31). Today's enterprises are using multiple types of experimental and analytical AI applications. One in ten enterprises now use ten or more AI applications. (Gartner).

One in ten enterprises now use ten or more AI applications. (Gartner)

Fig 31. Chatbots have displaced fraud detection as the top use of AI in 2019



Does your organisation use any of these artificial intelligence (AI) based applications? (2019: n = 2,791; 2018: n = 2,672; Multiple responses allowed.

Source: Gartner, 2019 Q3 Survey; COA: Hitec Analysis to the Inspiration of AI, Figure 1, 31 January 2019

The most popular AI use cases are:

- Chatbots (26% of enterprises)
- Process automation solutions (26%)
- Fraud analysis (21%)

Prevalent applications include:


- Consumer/market segmentation (15%)
- Computer-assisted diagnostics (14%)
- Call centre virtual assistants (12%)
- Sentiment analysis/opinion mining (12%)
- Face detection/recognition (11%)
- HR applications (e.g. CV screening) (10%)

Increasingly, certain applications are becoming widespread in particular industries.

- Nearly four in ten healthcare providers use computer-assisted diagnostics;
- Three in ten utilities use process automation tools;
- Six in ten healthcare payers, nearly half of financial service firms and four in ten insurers use AI for fraud detection;
- Three in ten retailers and a quarter of wholesalers use AI for consumer segmentation;
- A third of media companies use AI for sentiment analysis.

Natural language processing and computer vision AI underpin many of the popular and prevalent AI applications, including

Fig 32. Adoption of AI is uneven across, and within, sectors



chatbots, computer-assisted diagnostics, sentiment analysis and face detection. Companies are embracing AI's ability to replicate traditionally human activities in software for the first time – and the possibilities (including chatbots, computer-aided diagnostics and sentiment analysis) this enables.

Other, popular AI applications – fraud analysis, consumer segmentation and aspects of process automation – reflect AI's ability to identify patterns in data more effectively than traditional, rules-based software. As AI has expanded the breadth and complexity of workflows that can be automated, process automation has come of age. In 2017, given its potential, 64% of enterprises highlighted process automation as a focus for future AI deployment (Gartner). As solutions have matured, companies have made good on their intentions. In 2019, process automation is the joint most popular application for AI.

Sector adoption is in flux

Adoption of AI is uneven – across and within sectors – and in a state of flux. Sectors are diverging into 'early adopters' of AI, 'movers' and 'laggards'. Within sectors, adoption is dividing further among sub-sets of market participants.

Adoption of AI is uneven – across and within sectors – and in a state of flux.

Sectors are diverging into 'early adopters' of AI, 'movers' and 'laggards'.

'Early adopters' – sectors that proactively invested in AI – are reaping the benefits and maintaining their leadership. In 2017, financial services and high-tech & Telco companies anticipated increasing their investment in AI, in the following three years, more than companies in other sectors. Today, insurance, software & IT service and Telco companies lead in AI adoption (Fig. 32).

'Movers' have awoken to AI's potential and are closing the adoption gap. In 2017, adoption of AI in retail, healthcare and media was moderate relative to other sectors. Adoption in these sectors has accelerated. More than four in ten retail, healthcare and media companies have now invested in AI or will have done so within 12 months (Fig. 32).

'Movers' have awoken to AI's potential and are closing the adoption gap.

High rates of adoption in financial services, high-tech & Telco, retail, healthcare and media reflect the confluence of opportunity and engagement. AI offers extensive potential for value creation in these sectors. AI offer: numerous prediction and optimisation challenges well suited to AI; extensive data to train AI systems; quantifiable return on investment; and, to varying extents, the resources and ability to attract high-quality talent. Participants in the above sectors are also, typically, open to engaging with AI. 'Early adopters' meet opportunity with vision. 'Movers' have promptly recognised emerging opportunity – and begun to tackle impediments to adoption such as sprawling, siloed data estates.

'Laggards' – Government agencies, education companies and charities – are falling behind in AI adoption. While AI has potential to transform Government, in particular, given extensive data sets and numerous optimisation opportunities, AI engagement will continue to be inhibited by few AI initiatives, limited budgets for emerging technologies, siloed data and difficulty attracting AI talent. Individuals will engage with AI primarily as producers and consumers, not citizens, and in support of companies' and consumers' objectives. AI's transformation of western society will be led by companies, not governments, while vulnerable members of society will be among the last to benefit from AI.

Divergence is evident within as well as across sectors.

The proportion of insurance companies that have adopted AI, or intend to within the next 12 months, is ten percentage points higher than other financial service companies.

Within the healthcare sector, engagement with AI is greater among payers than providers. The value, and suitability, of particular AI use cases is driving 'hot spots' of activity within sectors. AI-powered fraud analysis, which can detect dishonest activity more effectively than traditional, rules-based systems, is the third most popular AI application today (Fig. 31) and is catalysing adoption among insurers and healthcare payers.

Interest in AI is diverging by department

A gulf is emerging between departments' interest in exploiting AI's potential. While IT departments express the greatest interest in AI, customer service teams are emergent AI champions (Fig. 33). The proportion of marketing, HR and finance departments interested in AI projects, meanwhile, is nearly double that of legal & compliance, sales and field service teams (Fig. 33).


A gulf is emerging between departments' interest in exploiting AI's potential.

Customer service teams' interest in AI reflects AI's value to both managers and workers, and low barriers to adoption. Customer service teams spend extensive time addressing repetitive, lower-value enquiries. AI, underpinned by natural language processing, enables replies to a growing proportion of enquiries to be created and sent automatically. For many other enquiries, contact centre workers' activities can be augmented through AI. Greater efficiency, and freedom to focus on higher-value cases, suits managers and workers alike. Tailwinds to engagement – including increasing adoption of contact centre software platforms, and the availability of 'best of breed' AI contact centre solutions such as DigitalGenius, in which we have invested – are fuelling interest.

Extensive interest in AI from marketers, similarly, reflects the breadth of marketing activities to which AI can be usefully applied and easily adopted. AI can augment customer segmentation, channel optimisation, content personalisation, price optimisation and churn prediction. Extensive training data is available and accessible for each activity, while uplift can be readily quantified.

Modest interest in AI from Legal & Compliance teams is at odds with AI's potential for value in these departments. While companies' legal and compliance costs are ballooning, AI powered by natural language processing can support activities including: automated time tracking; case law review; due diligence; litigation strategy; and communication compliance. Modest adoption of technology more broadly within legal departments, and cultural resistance to change, is inhibiting interest. Our primary research, however, highlights a divergence between innovative legal and compliance departments and laggards. Leaders are taking advantage of AI to gain significant competitive advantage. More broadly, we observe a tipping point in technology investment and openness to innovation among legal and compliance teams, as illustrated by the growth of "legal operations" personnel whose role is to optimise efficiency through modernisation and automation. Interest in AI among legal and compliance teams is likely to increase in the medium term.

Fig 33. A gulf is emerging between departments' interest in AI



Source: Gartner (June 2018)

AI leaders are better informed – and learning faster


Increasing AI adoption overall masks a growing gulf between leaders and laggards in AI – in companies' understanding, learning, strategy and investment.

Among AI laggards, fewer than two in ten believe they understand the technology¹, business, workplace – or industry implications of AI (Fig. 35, 'passives' and 'experimenters') (MIT Sloan Management Review). Among leaders ('pioneers' and 'investigators') the reverse is true: eight in ten understand its dynamics.

Laggards are set to fall further behind as their understanding of AI improves at a slower rate. In the last 12 months, between half and two thirds of AI leaders improved their understanding of AI to a great extent (Fig. 34) (MIT Sloan Management Review). During the same period, fewer than two in ten laggards did so.


Increasing AI adoption overall masks a growing gulf between leaders and laggards in AI – in companies' understanding, learning, strategy and investment.

Fig 34. The smart are getting smarter



Source: "Managing Business With Artificial Intelligence", MIT Sloan Management Review in collaboration with The Boston Consulting Group

Fig 35. There is a gulf between leaders' and laggards' understanding of the implications of AI.



Source: "Reshaping Business With Artificial Intelligence", MIT Sloan Management Review in collaboration with The Boston Consulting Group

Irrespective of their AI maturity, companies typically understand some considerations better than others (Fig. 35). Overall, companies are better attuned to the disruption AI will bring than the pragmatic challenges of deploying it. Companies understand best that: AI will change how companies generate value; that AI will shift industry power dynamics; and that an AI future will require different knowledge and skills to the past. Companies typically understand least: the costs of developing AI-based products and services; processes for algorithm training; and the effects AI will have on organisational behaviour.


Nine in ten AI pioneers – companies on the leading edge of AI deployment – have increased their investment in AI in the past year.

AI leaders are extending their advantage with greater investment

Companies proactively deploying AI are compounding their competitive advantage by increasing investment in AI at a greater pace than laggards.

Nine in ten AI pioneers – companies on the leading edge of AI deployment – have increased their investment in AI in the past year. Nearly two thirds companies investigating or experimenting with the technology have also done so. Among companies with no adoption or much understanding of AI, just one in five has increased spend on AI (Fig. 36, passives) (MIT Sloan Management Review).

Fig. 36. AI leaders are extending their advantage through greater investment



Source: J. Bontempo, P. Grosse, M. Reeves, D. Wilson, and H. Spier, "Artificial Intelligence in Business: Deep Dive," MIT Sloan Management Review and The Boston Consulting Group, September 2018.

Laggards are falling further behind in AI strategy

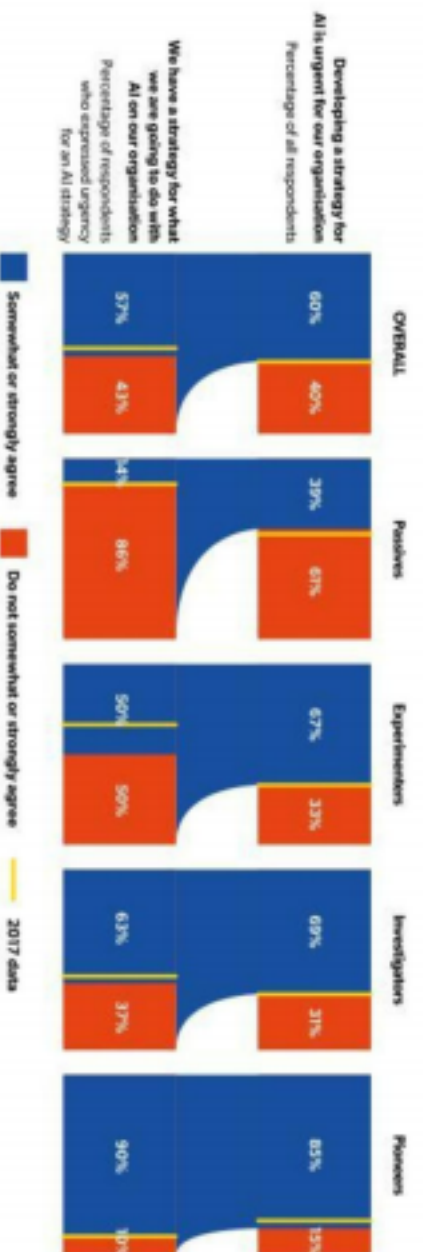
Laggards' sense of urgency regarding AI is weakening. The proportion of companies that believe developing an AI strategy is urgent for their organization is stable overall – at six in ten. However, while the proportion of proactive adopters with this belief has increased year-on-year, the proportion of laggards who share this view has fallen during the same period (Fig. 37, passives) (MIT Sloan Management Review).

Attitudes are shaping outcomes. Overall, the proportion of companies that have implemented an AI strategy has increased – but the proportion of laggards that have done so is unchanged (Fig. 37, passives) (MIT Sloan Management Review). AI leaders are compounding their advantages in understanding and learning with strategic planning – while laggards fall further behind.

The proportion of companies that have implemented an AI strategy has increased – but the proportion of laggards that have done so is unchanged.

(MIT Sloan Management Review)

Fig 37. While more companies have an AI strategy, the proportion of laggards with an AI strategy is unchanged




Source: S. Randelsham, P. Gerbert, M. Reeves, D. Kiron, and M. Spota, "Artificial Intelligence in Business Gets Real," MIT Sloan Management Review and The Boston Consulting Group, September 2018.

Leaders and laggards face different adoption challenges

The barriers to companies' adoption of AI are no longer consistent. Laggards are struggling with foundational considerations. They lack general technological capabilities to embrace AI, lack leadership support for AI initiatives, and are struggling to define use cases for the technology (Fig. 38, 'passives' and 'experimenters') (MIT Sloan Management Review).

Leaders' adoption challenges, in contrast, have shifted from 'if' to 'how'. Leaders have a strong understanding of AI use cases, extensive leadership support for AI initiatives and fewer technological constraints to AI adoption. Their challenges differ. Leaders are contending with the difficulties of attracting AI talent, balancing spend on AI with competing investment priorities and addressing cultural resistance to AI-led initiatives.

Fig 38. Leaders and laggards face different challenges to adoption



Source: S. Randelsham, P. Gerbert, M. Reeves, D. Kiron, and M. Spota, "Artificial Intelligence in Business Gets Real," MIT Sloan Management Review and The Boston Consulting Group, September 2018.


AI initiation has shifted from the C-suite to the IT department

Previously, the C-suite played a vital role in initiating AI projects, making technology decisions in relation to them, and approving project funding. Two years ago, Chief Executive Officers (CEOs), Chief Information Officers (CIOs) or Chief Technology Officers (CTOs) initiated two thirds of AI projects.

Today, just one in eight respondents highlight corporate leadership as the primary driver or initiator of AI projects. Interest in AI, and its initiation, has shifted from the C-suite primarily to the IT department (Fig. 39). The Customer Service function is also emerging as a powerful driver of AI projects.

AI engagement will continue to diffuse from the C-suite to lines of business. By providing ignition energy – identifying the disruptive potential of AI, prioritising experimentation with the technology and funding its deployment – the C-suite is necessary but insufficient to drive change. As companies' engagement with AI evolves from 'if' to 'how' – as understanding of AI use cases improves and implementation considerations weigh more heavily – line-of-business owners will play an ever-greater role in delivering value creation through AI.


Fig 39. Initiation of AI projects has shifted from the C-suite to the IT department



Companies prefer to buy, not build, AI

When adopting AI, more companies prefer to 'buy' than 'build'. Nearly half of companies favour buying AI solutions from third parties, while a third intend to build a custom solution internally (Fig. 40). Few companies – just one in ten – are prepared to wait for AI to be embedded in their favourite software products.

Fig 40. Nearly half of companies favour buying AI solutions from third parties



For many, a 'buy' strategy is appropriate given limited in-house AI capability and the proliferation of verticalised, best-of-breed software vendors with AI at the heart of their product propositions. In Europe alone 1,600 startups and scale-ups offer AI-led solutions, each focusing on a particular industry or business function (Chapter 7). Many offer best-in-class AI functionality, faster time to value and lower cost than developing in-house expertise and capability. Further, large buyers can frequently shape the product roadmaps of early stage companies to support their requirements. In sectors served by fewer early stage AI-led suppliers, such as Government and Education, propensity to 'build' is higher.

The low proportion of companies waiting for AI to be embedded in their favourite software products reflects buyers' urgency for AI and desire for sustainable competitive advantage. While democratising AI, incumbents are slower to embed AI features into existing solutions and less likely to offer best-in-class capability. By providing the same tooling to large groups of market participants, the competitive advantage they provide is also limited.

Paradigm shifts in technology typically destabilise incumbents and enthroned new winners. In 2019, as buyers prioritise capability and time to value, specialist suppliers are an attractive 'on-ramp' to AI. In time, as AI commodilises and buyers seek to consolidate and simplify their technology stacks, buyers may favour AI-enabled incumbents once again.

Workers are concerned about job security


Workers' views vary widely regarding the likely impact of AI on their daily activities – for example, whether AI will increase or decrease time spent with customers, or collaboration with colleagues (Fig. 41). As AI proliferates, on balance workers expect AI to increase the safety, quality and pace of their work while decreasing job security (Fig. 42).

Workers' expectations regarding the positive impact of AI on their roles are likely to be met. By augmenting existing workflows with new tools and capabilities, and increasing automation, quality of output and pace of productivity will increase.

Regarding workers' concerns about job security, AI is likely to enable the automation of select occupations that involve routine and repetition, such as telemarketing and truck

driving. In other roles, AI will augment workers' activities initially but displace a greater proportion of their activities over time – or obviate the need for additional hiring. In many cases, however, AI will simply augment and enrich individuals' roles, empowering workers with greater capabilities and the opportunity to focus on higher-value tasks. We discuss AI's potential to displace jobs, and other risks to society from AI, in Chapter 8.

Fig 41. Workers' views vary widely regarding the impact of AI on their activities



Sample Demographics

Sample size: n = 2,708, income greater than £10k, average age 41 years old
Education: Medium, Business (7%), Management: consumer, 25%,
high-level innovation, 10%.

Employees would prefer AI to be deployed as an enabler of higher, reducing routine work and riskier. Staff are very positive about how AI will impact their jobs, but they generally expect AI to increase safety and quality and decrease job security. They expect AI capabilities to increase safety, arrange scheduling the jobs it could replace in 30 years to more than four out of 10 jobs.

Source: Survey Analytics, How AI Will Impact Individuals From the Workers' Perspective, Quarter 2018

Fig 42. On balance, workers expect AI to decrease job security



Source: Survey Analytics, How AI Will Impact Individuals From the Workers' Perspective, Quarter 2018

Chapter 5



The advance of technology

Advances in AI technology are creating new possibilities. Custom silicon is enabling a new generation of AI hardware. Emerging software techniques are delivering breakthroughs in multiple domains and decoupling progress from the constraints of human experience.

Summary

- While graphical processing units (GPUs) catalysed AI development in the past, and will continue to evolve, hardware innovations are expanding AI's potential. Hardware is being optimised, customised or re-imagined to deliver a new generation of AI accelerators.
- Hardware with 'tensor architectures' is accelerating deep learning AI. Vendors, including NVIDIA and Google, are optimising or customising hardware to support the use of popular deep learning frameworks.
- We are entering the post-GPU era. Leading hardware manufacturers are creating new classes of computer processor designed, from inception, for AI. Custom silicon offers transformational performance and greater versatility.
- Custom silicon is also taking AI to the 'edge' of the internet - to IoT devices, sensors and vehicles. New processors engineered for edge computing combine high performance with low power consumption and small size.
- As quantum computing matures, it will create profound opportunities for progress in AI and enable humanity to address previously intractable problems, from personalised medicine to climate change. While nascent, quantum computing is advancing rapidly. Researchers have developed functioning neural networks on quantum computers.
- Reinforcement learning (RL) is an alternative approach to developing AI that enables a problem to be solved without knowledge of the domain, instead of learning from training data. RL systems reward and reinforce progress towards a specified goal. AlphaGo Zero, an RL system developed by DeepMind to play the board game Go, developed unrivalled ability after just 40 days of operation. In 2019, developments in RL will enable groups of agents to interact and collaborate effectively.
- Progress in RL is significant because it decouples system improvement from the constraints of human knowledge. RL is well suited to creating agents that perform autonomously in environments for which we lack training data.
- Transfer learning (TL) enables programmers to apply elements learned from previous challenges to related problems. TL can deliver stronger initial performance, more rapid improvement and better long-term results. Interest in TL has grown seven-fold in 24 months and is enabling a new generation of systems with greater adaptability.
- By learning fundamental properties of language, TL-powered models are improving the state of the art in language processing - in areas of universal utility. 2018 was a breakthrough year for the application of TL to language processing.
- TL is also enabling the development of complex systems that can interact with the real world, delivering systems with greater adaptability, and supporting progress towards artificial general intelligence, which remains far from possible with current AI technology.
- Generative Adversarial Networks (GANs) will reshape content creation, media and society. An emerging AI software technique, GANs enable the creation of artificial media, including pictures and video, with exceptional fidelity. GANs will deliver transformational benefits in sectors including media and entertainment, while presenting profound challenges to societies - beware, fake news 2.0.

Recommendations

Executives

- Ensure your organisation, or suppliers, are taking advantage of the latest advances in AI hardware for faster solutions to more complex challenges.
- Custom silicon for edge computing is enabling ‘edge’ devices – drones, robots, embedded devices and sensors – with greater AI capabilities. Explore whether AI-enabled edge applications could offer your company, or customers, utility.
- Reinforcement learning can be usefully applied to tackle problems of control (such as warehouse automation) and coordination (including fleet optimisation). Explore whether reinforcement learning could deliver efficiency improvements and cost savings for your company.

Entrepreneurs

- Take advantage of hardware with tensor architectures to accelerate the development of deep learning systems.
- Offer more advanced language processing in your solutions by drawing on recent breakthroughs in transfer learning.
- Generative Adversarial Networks (GANs) can be usefully applied to a wide variety of domains beyond media, from signal normalisation to network security. Explore whether they could provide utility for your application.

Investors

- The ‘post-GPU era’ will create new winners. Explore companies developing custom silicon for AI, for the data centre and edge devices.
- Reinforcement learning offers solutions to a range of challenging problems. Identify companies taking advantage of reinforcement learning for competitive advantage.
- Identify opportunities for portfolio companies to take advantage of advances in computer vision and language enabled by transfer learning.
- Explore the field of quantum computing. While nascent, it will gain significance rapidly in the years ahead.

Policy-makers

- Transfer learning, reinforcement learning and generative AI enable AI systems with greater capability and adaptability – and pose new risks to society. Explore the implications of emerging AI technology in Chapter 8.
- The UK is an emerging leader in the nascent field of quantum computing. Review the National Quantum Technologies Programme to explore the UK’s strengths and challenges in quantum technology and identify opportunities for policy-makers’ support.

Explore our AI Playbook, a blueprint for developing and deploying AI, at www.mmcventures.com/research.

AI hardware is being optimised, customised and reimaged

Training the neural networks that power many AI systems is computationally intensive. Graphical Processing Units (GPUs) – hardware that is efficient at performing the matrix mathematics required – have enabled extensive progress and transformed the field of AI (see Chapter 3). In the last decade, computing performance for AI has improved at a rate of 2.5x per year (IBM). The performance of GPUs will continue to increase.

However, GPUs were designed for graphics processing – not AI. Manufacturers exploited GPUs' ability to perform matrix calculations when it became apparent that AI benefited from the same mathematics. Frequently, just a third of a GPU's core area is used for AI.

As AI matures, greater demands are being placed on the hardware that powers it. Larger data sets, 'more model parameters, deeper networks, moving AI to 'edge' devices, and an ambition to tackle new challenges demand improved capability. "Current hardware was holding developers back," (Nigel Toon, Graphcore)

Below, before describing breakthroughs in AI software techniques, we highlight three dynamics shaping AI hardware – the optimisation, customisation and reimagination of hardware for AI.

- Competition among hardware providers is fierce. In response to recent industry benchmarking, which compared Google's and NVIDIA's processors (<https://mlperf.org/results/1>), both parties claimed victory (<https://bit.ly/2l6wMzT>, <https://bit.ly/25YLEQd>). Developers and consumers alike will benefit from intense competition, as new hardware:
- lowers the cost of compute for AI, democratising access and accelerating proliferation of the technology;
 - increases the speed at which systems can be trained and iterated, shortening development cycles;
 - reduces required power consumption, enabling AI on 'edge' devices such as Internet of Things (IoT) units, autonomous vehicles, implanted medical devices and sensors; and
 - enables more complex and effective models.
- Better models can improve existing applications and

enable new ones (in December 2018, Google used sophisticated deep learning to predict the 3D structure of proteins, based solely on their genetic sequences, for the first time).

- accelerates new approaches to AI, such as reinforcement learning (RL) and transfer learning (TL), which we explain below.

Tensor architectures are accelerating deep learning

Deep learning AI continues to offer myriad breakthroughs and benefits – in domains including computer vision and language and applications ranging from autonomous vehicles to medical diagnosis and language translation.

In response, vendors are optimising or customising hardware to support the use of popular deep learning frameworks. While addressing a more limited set of instructions, this hardware enables faster system training and performance from common AI frameworks – with varying degrees of specificity.

Vendors are optimising or customising hardware to support the use of popular deep learning frameworks.

NVIDIA has introduced GPUs with architectures optimised for deep learning on a range of frameworks. The Company's Tesla GPUs contain hundreds of Tensor Cores that accelerate the matrix calculations at the heart of deep learning AI. Tesla GPUs deliver faster results with common AI frameworks, particularly convolutional neural networks used for computer vision systems.

Tesla GPUs enable suitable neural networks to be trained in a third of the time previously required (Fig. 4-3) and operate four times faster (Fig. 4-4). Compared with a traditional CPU, Tesla GPUs offer a 27-fold improvement.

Fig 43. Tesla GPUs enable suitable neural networks to be trained in a third of the previous time





Fig 44. Tesla GPUs allow suitable neural networks to operate four times faster than previously



Google's Tensor Processing Unit (TPU) is an application-specific integrated circuit (ASIC) – a custom microchip – designed specifically to accelerate AI workloads on the popular TensorFlow framework.

After publicising its use of TPUs in May 2016, Google announced its second-generation TPU in May 2017 and third generation in May 2018. While first-generation TPUs were limited to inferencing (processing queries through a trained network), subsequent generations accelerate system training as well as inference.

Optimised to process the mathematics required by TensorFlow, TPUs offer exceptional performance for TensorFlow applications. Even moving from Google's second-generation TPU to its third reduced by nearly 40% the time required to train ResNet-50, an industry-standard image classification model.

Fig 45. Google's second-generation TPU reduced the time required to train an image



Initially, Google used TPUs only within its own data centres, to accelerate Google services including Google Photos (one TPU can process 100 million photos per day). Google Street View and Google's RankBrain search facility. TPUs are now accessible to general developers and researchers via the Google Cloud Platform.

The post-GPU era: custom silicon is enabling new possibilities

Leading hardware manufacturers are diverging from architectures used in the past. In 2019 a new class of computer processors designed, from inception, for AI will emerge. Custom silicon, designed from first principles for AI, offers transformational performance, capability similar to existing systems for a fraction of the power or space, and greater versatility.

Incumbent microchip manufacturers, global technology companies and dozens of disruptive early stage companies, including Cerebras, Graphcore and Mythic are developing next-generation processors for AI.

Fig 46. Graphcore's IPU is designed, from inception, for AI




Source: Graphcore

Graphcore, a privately-held 'scale-up' company in the UK that has attracted over \$300m of venture capital funding, has developed an Intelligence Processing Unit (IPU) (Fig. 46). Graphcore's IPU combines a bespoke, parallel architecture with custom software to offer greater performance than existing systems. Graphcore's benchmarking suggests that its IPU can deliver 200-fold performance improvements in selected tasks, compared with GPUs (Fig. 47).

The IPU's architecture and software enable large quantities of data to be consumed in parallel, instead of sequentially, and from multiple locations ('graph computing' in place of 'linear addressing'). Data is transported across the IPU's 1,000+ sub-processors more efficiently, and the IPU provides faster access to greater volumes of memory to reduce bandwidth limitations.

As well as enabling existing workloads to be processed more rapidly, new hardware architectures such as IPUs will enable developers to tackle previously intractable challenges.

Fig 47. Graphcore's IPU could deliver 200-fold performance improvements in selected tasks



Source: Graphcore

Custom silicon is taking AI to the edge

While cloud computing proliferates, a 'barbell' effect is emerging as a new class of AI hardware is optimised for edge computing instead of the data centre.

Edge computing moves the processing of data from the cloud to the 'edge' of the internet – on to devices where it was created such as autonomous vehicles, drones, sensors and IoT devices. Increasingly, edge computing is required – as devices proliferate, and connectivity and latency issues demand on-device processing for many.

Numerous hardware manufacturers are developing custom silicon for AI at the edge. In October 2018, Google released Edge TPU – a custom processor to run TensorFlow Lite models on edge devices. A plethora of early stage companies, including Gyrfalcon, Mythic and Syntiant are also developing custom silicon for the edge.

In 2019, as well as enabling next generation AI in the cloud, custom silicon will transform AI at the edge by coupling high performance with low power consumption and small size.

Quantum computing will unlock profound opportunities

Quantum computing is a paradigm shift in computing that exploits the properties of nature – quantum mechanics – to offer profound new possibilities. While nascent, quantum computing hardware and software are advancing rapidly. 2019 may be the year of 'quantum supremacy' – the first time a quantum computer solves a problem a classical computer cannot.

Quantum hardware, and associated software to accelerate AI, are emerging. In addition to building quantum processors, Google is developing quantum neural networks. In November 2018, an Italian team of researchers developed a functioning quantum neural network on an IBM quantum computer (<https://bit.ly/2GxIpeel>). Rigetti, a manufacturer of quantum computers and software, has developed a method for quantum computers to run certain AI algorithms.

While quantum computing technology will take time to mature, in the decade ahead quantum-powered AI will enable humanity to address previously intractable problems – from climate change to personalised medicine.

In 2019, as well as enabling next generation AI in the cloud, custom silicon will transform AI at the edge by coupling high performance with low power consumption and small size.

Breakthroughs in software development are delivering transformational results

While novel hardware will enable more powerful AI, recent breakthroughs in software development are delivering transformational results.

Below, we explain how advances in two alternative approaches to developing AI systems – RL and TL – are enabling the creation of programs with unrivalled capabilities. We also describe how a new AI software technique – the Generative Adversarial Networks (GAN) – has reached a tipping point in capability that will reshape media and society.

Reinforcement learning is creating powerful AI agents

Recent advances in RL, an alternative approach to developing AI systems, are delivering breakthrough results – and raising expectations regarding AI's long-term potential.

Typically, an AI system analyses training data and develops a 'function' – a way of relating an output to an input – that is used to assess new samples provided to the system (supervised learning).

RL is an alternative approach that uses principles of exploration and reward. Human parents encourage children's development through emotional rewards (limiting, clapping and verbal encouragement) and physical prizes (toys or sweets). Similarly, after an RL system is presented with a goal, it experiments through trial and error and is rewarded for progress towards the goal. While the system will initially have no knowledge of the correct steps to take, through cycles of exploration RL systems can rapidly improve.

RL is an efficient approach for teaching an agent to interact with its environment. Developers begin by specifying a goal and elements within the agent's control – for example, in robotics, the joints that a robot can move and the directions in which it can travel. By rewarding useful progress and negatively reinforcing failure, as early as 1997 it was demonstrated that RL could produce a robot that walked in a dynamic environment – without knowledge of the environment or how to walk (Benbrahim and Franklin).

Developments in RL are enabling profound milestones in the training of individual AI agents and, by teaching cooperation, groups.

18 months ago AlphaGo Zero, an RL system developed by DeepMind to play the board game Go, outperformed DeepMind's previous AI Go system that had been trained


using traditional, supervised learning. Provided only with the rules of Go, and without knowledge of any prior games, by playing against itself AlphaGo Zero reached the level of AlphaGo Master in 21 days. After 40 days, AlphaGo Zero surpassed all prior versions of AlphaGo to become, arguably, the strongest Go player in the world (Fig. 48). "Humans seem redundant in front of its self-improvement" (Ko Je, World No. 1 Go player).

15 months ago, DeepMind developed a more general program – AlphaZero – that could play Chess, Shogi and Go at levels surpassing existing programs.

RL is well suited to creating agents that can perform autonomously in environments for which we lack training data, and enabling agents to adapt to dynamic environments. In 2019 RL will catalyse the development of autonomous vehicles. In the longer term the exploration of space, where training data is limited and real-time adaptation is required, is likely to draw on RL.

Progress in RL is significant, more broadly, because it decouples system improvement from the constraints of human knowledge. RL enables researchers to "achieve superhuman performance in the most challenging domains with no human input" (DeepMind). We explore this profound implication of AI in Chapter 8.

Fig 48. Reinforcement learning enabled AlphaGo Zero, a system developed by DeepMind to play Go, to achieve unrivalled capability after 40 days of play.



Reinforcement learning is enabling multi-agent collaboration

In 2019, developments in RL will also enable groups of agents to interact and collaborate with each other more effectively.

Games, which present a safe and bounded environment for learning, are valuable for training RL systems (Aditya Kaul). Defence of The Ancients 2 (Dota2) is a cooperative online game for teams of five players (Fig. 49). While previous environments required AI agents to optimise only for their own success when responding to the actions of other teams, Dota2 requires agents to consider the success of their team.

OpenAI 5 is a Dota2 team developed by OpenAI, a non-profit AI research company building safe artificial general intelligence. OpenAI used RL in a similar manner to DeepMind's AlphaGo Zero to train its team.


OpenAI 5 agents initially played against themselves to learn individual and cooperative skills. Subsequently, they were able to improve rapidly (Fig. 50) and defeat all but the top professional human teams.

Developing RL remains challenging. Designing reward functions can be difficult as RL agents will 'game the system' to obtain the greatest reward. OpenAI discovered that if they offered agents rewards for collecting power-ups, which would enable the agents to complete their task faster, agents abandoned the task to collect the power-ups given the available rewards. Even with sound reward functions, it can be difficult to avoid 'overfitting' solutions to their local environment.

Fig 49. Reinforcement learning is enabling effective multi-agent collaboration (AI agents playing Defence of the Ancients 2)



Fig 50. Reinforcement learning enabled the OpentAI 5 team to surpass rapidly the performance of most human teams




Source: OpentAI

Transfer learning is delivering breakthroughs in language AI – and beyond

Traditional AI requires systems to be trained from a standing start, which demands data and time, or accepting the outputs of existing, pre-trained networks whose training data is inaccessible. Accordingly, AI development is frequently inefficient or sub-optimal.

Transfer learning (TL) is an emerging approach for developing AI software, which enables programmers to create novel solutions by re-using structures or features of pre-trained networks with their own data. By drawing upon skills learned from a previous problem, and applying them to a different but related challenge, TL can deliver systems with stronger initial performance, more rapid system improvement, and better long-term results (Fig. 51).


Fig 51. Transfer learning can offer strong initial performance, faster improvement and better long-term results



Source: Tony Shalun

TL has been used to accelerate the development of AI computer vision systems for over a decade. In the last 24 months, however, interest in TL has grown 7-fold (Fig. 52). In 2019 TL is being applied to broader domains – particularly natural language processing.

Fig 52. Interest in transfer learning has grown 7-fold in 24 months



Source: Google Trends

To date, natural language processing has operated at a shallow level, struggling to infer meaning at the level of sentences and paragraphs instead of words. Word embedding, an historically popular technique for inferring the meaning of a word based on the words that frequently appear near to it, is limited and susceptible to bias. The absence of extensive, labelled training data for natural language processing has compounded practitioners' challenges.

By enabling better results with less training data, transfer learning is delivering transformational results. 2018 was a breakthrough year for the application of transfer learning in language processing.

By enabling better results with less training data, TL is offering transformational results. 2018 was a breakthrough year for the application of transfer learning in language processing.

- In March 2018, the Allen Institute for Artificial Intelligence used TL to deliver ELMo (Embeddings from Language Models), which improved the state of the art for a broad range of natural language tasks including question answering and sentiment analysis (<https://bit.ly/2HY6IMZ1>).
- In May 2018, research institution Fast AI released ULMFiT (Universal Language Model Fine-tuning for Text classification). ULMFiT underscored that TL can be applied to language processing tasks and introduced techniques for fine-tuning language models. By using TL, with only 100 labelled examples ULMFiT matched the performance of systems trained with 100-fold more data. Their method also offered improved text classification and reduced error rates by 18-24% on many data sets (<https://bit.ly/2Hmu1d1>).
- In mid-2018, OpenAI demonstrated the ability to achieve impressive results on a diverse range of language tasks from a single starting point. OpenAI's general, task-agnostic model outperformed models that used architectures specifically crafted for tasks including question answering and textual entailment (<https://bit.ly/2t9GjM1>).

- In October 2018, Google open-sourced BERT (Bidirectional Encoder Representations from Transformers), an RL-based language processor that achieved state of the art results on 11 natural language processing benchmarks (<https://bit.ly/2QqgmY30J>). The "bidirectionality" of BERT allows context to be carried between sentences, for improved textual responses.

New, TL-powered models "learn fundamental properties of language" (Matthew Peters, ELMo). By doing so, they may unlock higher-level capabilities in language processing with universal utility – including text classification, summation, text generation, question answering and sentiment analysis.

Transfer learning enables complex systems to interact with the real world

In many situations, gathering data to train AI systems is laborious, expensive or dangerous. Amassing data to train an autonomous vehicle, for example, could require millions of hours of labour, billions of dollars and considerable risk. Simulation, combined with transfer learning, offers a solution. Instead of capturing real-life data, environments are simulated. Using TL, learnings from the simulation can then be applied to the real-world asset.

In the field of robotics, similarly, training models on real-world robots is slow and costly. Learning from a simulation, and transferring the knowledge to a physical machine, can be preferable.

TL may be "a pre-requisite for large-scale machine learning projects that need to interact with the real world" (Sebastian Ruder). As a result, "transfer learning will be the next driver of machine learning commercial success" (Andrew Ng).

Transfer learning offers adaptability and progress towards artificial general intelligence

TL offers profound as well as pragmatic benefits.

By reducing the volume of training data required to solve a problem, TL enables humans to develop systems in domains where we lack large numbers of labelled data points for system training.

By offering greater adaptability, TL also supports progress towards artificial general intelligence (AGI) – systems that can undertake any intellectual tasks a human can perform. While AGI is far from possible with current AI technology, developments in TL are enabling progress. "I think transfer learning is the key to general intelligence. And I think the key to doing transfer learning will be the acquisition of conceptual knowledge – knowledge that is abstracted away from perceptual details of where you learned it, so you can apply it to a new domain" (Dennis Hassabis, DeepMind).

"I think transfer learning is the key to general intelligence."

Dennis Hassabis, DeepMind

GANs will transform media and society

First proposed in 2014, Generative Adversarial Networks (GANs) are a novel, emerging AI software technique for the creation of lifelike media – including pictures, video, music and text. Exceptional recent progress in the development of GANs (Fig. 53) has enabled breakthrough results. Today, GANs can generate highly realistic media, which – despite being artificially generated – are virtually impossible to differentiate from real content (Fig. 54)


Today, GANs can generate highly realistic media, which – despite being artificially generated – are virtually impossible to differentiate from real content.

Fig 53. GANs' ability to create lifelike media has rapidly improved



Source: Goodfellow et al, Radford et al, Liu and Tsumi, Karras et al, <https://bit.ly/2GwTl6t>

Fig 54. GANs can generate artificial images that appear real (none of these individuals exist)



While GANs are frequently used to create images, their utility is broader. Additional uses include:

- **Alternative media:** GANs can create different forms of media, such as music or text in the style of particular individuals.
- **System training:** GANs can be used to improve the training of AI classification systems. Neural networks used for image classification are easily misled by minor changes to images, including those invisible to the human eye. A classifier can be made more robust by using it as a GAN discriminator, and using the GAN to create altered images.
- **Data manipulation:** Frequently, it is important to remove personal information from data – such as the number plate of a vehicle or the face of a child in an image. Combining GANs with additional techniques, such as autoencoding, enables the addition or removal of features from data.
- **Data normalisation:** GANs enable data from different sources to be normalised. Instead of feeding random noise into a GAN's generator, developers can input types of signal data that are different from the desired output. The GAN will normalise the result. For example, health data collected from different devices will have different sampling frequencies and accuracy tolerances. GANs can normalise the signals for greater comparability.
- **Network security:** Because GANs are structured to distinguish between the real and the counterfeit, they are valuable for domains such as cybersecurity where it is a priority to detect anomalies in network access or activity.
- **Data creation:** AI classification systems are inhibited by the volume and quality of data available to train them. GANs can produce additional training data to improve classifiers' accuracy. This technique has been used to improve the classification of liver lesions. Creating data using GANs poses challenges as well as opportunities. The GAN discriminator will have been trained using a limited data set. While the generator's outputs may appear realistic, the images produced may not correctly reflect the appearance of a human body with the same disease.

GANs will deliver transformational benefits. The ability to generate lifelike images to a desired specification will reshape the media sector. Further, GANs will enable agencies to capture footage of brand ambassadors and then repurpose footage to create an infinite range of convincing variations. Ambassadors could appear to speak in foreign languages (to promote goods and services in international markets) and discuss new products – without recording any additional footage.

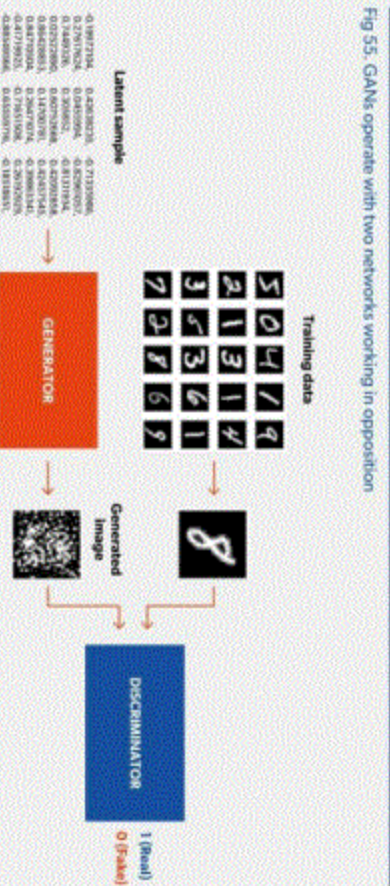
GANs also present profound ethical and societal risks. GANs can be used to: splice individuals' faces onto existing video without their consent; develop video in which individuals appear to speak words they have not spoken; create counterfeit evidence for criminal cases; and generate or alter footage to create 'fake news'. We discuss the implications of GANs for society in Chapter 8.

GANs will deliver transformational benefits. They also present profound risks. We discuss the implications of GANs in Chapter 8.

GANs operate by two networks working in opposition

GANs operate by two networks – a ‘generator’ and ‘discriminator’ – working in opposition to create increasingly lifelike media.

For a visual GAN, a generator receives a random input, such as a matrix of numbers, and follows a series of mathematical transformations to convert the input into a picture. Initial results will be poor, resembling random sets of pixels (Fig. 55).



Source: Nick Szegedy

GANs operate by two networks – a ‘generator’ and ‘discriminator’ – working in opposition to create increasingly lifelike media.

The output of the generator is then passed to the discriminator. The discriminator is a separate convolutional neural network that has been trained to recognise counterfeit images of the type in question – in this example, handwritten digits. The discriminator assesses whether the image received from the generator is authentic or has been artificially generated. Following the discriminator's decision, the correct answer is revealed.

If the discriminator correctly determines that the output is artificially generated, the generator changes the weights in the network responsible for the output recognised as counterfeit, and reinforces the weights in the discriminator that led to the correct conclusion.

If the discriminator incorrectly assesses the output from the generator, the weights in the generator, which led to a useful image, are reinforced, and the features in the discriminator, which led to an incorrect result, are down-weighted to yield a better assessment in future.

As the two networks work in parallel, influencing one another, the output from the generator improves until the accuracy of the discriminator is no better than chance (a 50/50 probability of correctly determining the authenticity of the generated image).

**The discriminator assesses
whether the image received
from the generator is
authentic or has been
artificially generated.**

Chapter 6



The war for talent

While demand for AI professionals exceeds supply, winners and losers are emerging in the war for talent.

Summary

- Demand for AI talent has doubled in 24 months. There is a gulf between demand and supply, with two roles available for every AI professional.
- The pool of AI talent remains small. AI demands advanced competencies in mathematics, statistics and programming. AI developers are seven times more likely to have a Doctoral degree than other developers.
- Supply is increasing – machine learning has become the top emerging field of employment in the United States. Greater supply is being driven by: high pay; the inclusion of AI modules in university computer science courses; companies' investment in staff training; and AI technology companies 'pump priming' the market with free educational resources.
- Over time, AI tools offering greater abstraction will make AI accessible to less specialised developers.
- Talent shortages are sustaining high salaries. AI professionals are among the best paid developers and their salaries continue to increase: half enjoyed salary growth of 20% or more in the last three years.
- Winners and losers are emerging in the war for talent. The technology and financial services sectors are absorbing 60% of AI talent.
- The 'brain drain' from academia to industry is real and will have mixed implications, catalysing AI's immediate impact while inhibiting teaching and moving value from the public domain to private companies.
- High job satisfaction is intensifying the war for talent. Three quarters of AI professionals are satisfied in their current role.
- To optimise hiring and retention, companies should align roles to AI professionals' primary motivators – learning opportunities, office environment and access to preferred technologies.
- New practitioners in the field are following sub-optimal paths to employment. Company websites and technology job boards are less effective than engaging with recruiters, friends, family and colleagues, according to those already employed in the field.

Recommendations

Executives

- To attract AI talent, leverage your advantages as a large company. Offer access to vast data sets, the opportunity for impact at scale and high salaries.
- Develop best-in-class training to up-skill existing developers.
- Diversity delivers economic value and competitive advantage. Review the culture in your company, AI team and hiring practices to ensure diversity, representation and inclusion.
- Collaborate with universities to support your search for talent, strengthen your reputation as a supporter of AI innovation and train colleagues through engagement with university research programmes.

Entrepreneurs

- Engage with universities, meet-ups and conferences to identify and attract promising candidates before they enter the market.
- Exploit your advantages as a start-up to attract AI talent. Offer work that can 'make a difference', direct impact on product, opportunities for learning, access to preferred technologies and an appealing office environment – in addition to equity.
- Follow best practices in our AI Playbook (www.mmventures.com/research) to optimise each stage of your recruitment funnel.
- Given demand for AI talent, maintain a focus on team satisfaction and retention.

Investors

- Amidst a competitive market for talent, assess potential investees' ability to attract and retain high quality AI personnel.
- Develop a competency in the recruitment of AI talent, including engagement with specialist recruiters, to assist portfolio companies.
- Understand best practices for every stage of a company's AI recruitment funnel – and ensure their proliferation across your portfolio.

Policy-makers

- Develop programs and funding to support education in science, technology, engineering and mathematics (STEM) subjects. Investment in STEM will mitigate talent shortages and empower workers for the age of AI.
- Enable the next generation of AI academics and mitigate the 'brain drain' to industry by providing greater, more accessible grant funding and access to national data sets for the public good.

Explore our AI Playbook, a blueprint for developing and deploying AI, at www.mmventures.com/research.


AI talent is in high demand

As AI is woven into the fabric of consumer experiences, and corporate adoption of AI extends from early adopters to the early mainstream, demand for developers who can create AI solutions has surged. In the United Kingdom, job listings for AI roles have increased 465% since 2014 (indeed). A quarter of companies highlight that lack of available AI talent is a primary inhibitor in their efforts to adopt AI (Gartner).

In the last 24 months, AI-related job postings as a proportion of total postings nearly doubled.

Growth in demand is accelerating. In the United States, year-on-year growth in AI-related job postings increased from 20% (2016) to 32% (2017) (indeed). In the last 24 months, AI-related job postings as a proportion of total postings nearly doubled (Fig. 56).


Fig 56. AI-related job postings as a proportion of total job postings have doubled in 18 months



Supply is increasing...

In the United States, machine learning has become the top emerging field of employment, with ten times the number of individuals listing it as their profession today compared with five years ago (LinkedIn) (Fig. 57). Data science, more broadly, is the second-from-top emerging field of employment, with more than six-fold growth.

Fig 57. In the US, machine learning is the top emerging job



Source: LinkedIn

Supply is increasing as:

- developers recognise opportunity for challenge and high pay within the field;
- universities update computer science courses to include AI components and resources. Undergraduate Computer Science courses at the universities of Cambridge, Harvard, MIT, Oxford, Princeton and Stanford all include AI components. In addition, many universities offer free online AI resources, including Stanford's and Columbia's 'Machine Learning' courses and MIT's 'Deep Learning for Self-Driving Cars' (Fig. 58).




MIT's 'Deep Learning for Self-Driving Cars'

Fig 58. Many universities offer free online AI courses and resources

Increase in individuals' listed professions compared with five years ago

- large companies invest in training initiatives for staff. Three quarters of large companies are offering some form of in-house or external training program, with a third providing formal training (Gartner).
- AI-focused technology companies 'pump prime' the market with free educational resources (Fig. 59), including Google's 'Machine Learning' course and NVIDIA's 'Fundamentals of Deep Learning for Computer Vision' resource.




'Deep Learning By Google' course

Fig 59. Technology companies are offering free educational AI resources

...but the pool of AI talent is small


Estimates of the number of global AI developers vary widely, in part depending upon definition. There may be as few as 22,000 highly-trained AI specialists (Element) and up to 300,000 AI researchers and practitioners within broader technical teams (Fercenti). AI originated in academia. The advanced mathematics, statistics and computer science required to understand and apply AI required extensive education, limiting the size of the available talent pool. AI developers are highly educated; nearly 60% have a Master's or Doctoral degree (Fig. 60). AI developers are twice as likely to have a Master's degree and seven times more likely to have a Doctoral degree than other professional developers (Fig. 61). Two thirds of data scientists believe their career success (Kaggle).

Fig 60. 60% of AI developers have a Master's or Doctoral degree.



Source: Kaggle

Fig 61. AI developers are seven times more likely to have a Doctoral degree than others



Source: Kaggle, Stackoverflow

In addition to technical skills, increasingly AI practitioners must have:

- domain knowledge, to interpret data appropriately and provide relevant recommendations;
- engineering experience, to develop solutions that work in the real world as well as the laboratory;
- commercial experience, to develop and manage AI teams.

The combination of technical, sector-specific, engineering and commercial competencies required from AI professionals continues to limit the size of the talent pool.

Education and the democratisation of AI will mitigate talent shortages

Over time, a larger talent pool and more accessible AI tools will alleviate much of the shortfall in AI talent – and enable greater realisation of AI's benefits.

Governments' investment in education – in science, technology, engineering and mathematics (STEM) subjects – will be vital for countries to broaden their pools of AI talent. The proliferation of AI courses and resources from universities and technology companies, and market demand, will also boost supply.

However, AI will also become accessible to less specialised developers over time. Development environments for new technologies tend towards higher levels of abstraction over time (new developers program in assembly language today). AI will follow this pattern.

Prior to 2000, developing AI required advanced mathematics, sophisticated programming and the specification of algorithms by hand. Successive developments have reduced the burden on developers:


- NumPy (2005) abstracted portions of required mathematics;
- Compute Unified Device Architecture (CUDA, 2007) reduced the requirement to code by hand;
- Python libraries (2010) and TensorFlow (2015) progressively abstracted network development.

Today, Google, Amazon and Microsoft offer AI services that require no implementation knowledge of AI. Developers with limited coding experience can upload data and solve simpler classification problems. While there will remain a large core of highly educated AI developers to progress research, advanced- and domain-specific AI, we expect the technology to become accessible to a greater proportion of developers over time, expanding the pool of developers who can deploy it.

There is a gulf between demand for AI talent and supply

While supply of AI talent is increasing, demand significantly outstrips supply and will continue to do so in the medium term. There are 2.3 roles available for every suitable candidate (indeed, “There is a mountain of demand and a trickle of supply” (Chris Nicholson, CEO, Splynn)). AI professionals themselves cite the lack of available talent as their second-greatest challenge (Fig. 62).

Fig 62. AI professionals cite the lack of available talent as their second-greatest challenge



Barriers faced at work by data scientists

Source: Splynn

Talent shortages are sustaining high salaries

A shortage of AI developers is driving high salaries in the market. Data scientists and machine learning specialists are among the best paid professional developers (Fig. 63). At the 20 highest-paying companies, salaries for AI engineers average \$224,000 (Fig. 64). Leaders in the field command vast sums.

Fig 6.3. AI professionals are among the highest paid developers






Fig 6.4. Salaries for AI engineers average \$224,000 at the 20 highest-paying companies



AI developers' salaries are particularly high relative to their level of professional experience. Nearly half of data scientists have under two years of professional experience (Kaggle); nearly three quarters have less than ten. Compared with other developers, data scientists enjoy among the greatest salary premium relative to their level of experience (Fig. 65).

Fig. 65. AI professionals are paid highly relative to their level of experience




AI salaries continue to increase

Salaries for AI professionals have grown significantly in recent years and continue to increase. Almost all data scientists report increased pay in the last three years; nearly half grew their salaries by 20% or more (Fig. 66).

In the last 12 months, salaries have continued to increase (Fig. 67). This year's pay dynamic has been more favourable to AI professionals than to many other developers. However, AI professionals are not the only group to enjoy significant year-on-year pay rises. Developers specialising in system administration, embedded applications and enterprise applications all received similar increases. DevOps specialists, who integrate and automate development and operations functions for faster cycles of improvement, are enjoying the greatest average raises.


Fig 66. Most AI professionals' salaries have increased in the last three years



Source: PwC

Salaries for AI professionals have grown significantly in recent years and continue to increase.

Fig 67. AI professionals' salaries have increased further in the last 12 months



Source: Statista/OWS, MWC, Workday

Winners and losers are emerging in the war for AI talent

Despite AI's potential to reshape sectors ranging from retail to healthcare, technology and financial services firms are absorbing nearly 60% of AI talent (Fig. 68) (Burch Works). 44% of data scientists are employed in the Technology sector – more than in the healthcare, consulting, marketing, retail, academia and Government sectors combined. Financial services, with a 14% share of data scientists, is a distant second.

Within the Technology sector, the world's largest technology companies – including Amazon, Apple, Facebook, Google, IBM and Microsoft – are consolidating much of the available talent. Amazon, Microsoft and Apple combined are estimated to be investing \$620m in AI talent (PwC).



Fig 68. Technology and financial services firms are absorbing nearly 60% of AI talent (distribution of data scientists by sector)



Source: The Burch Works Study – May 2018, N=2,212

The technology and financial services sectors are emerging winners in the war for AI talent – and creating virtuous cycles to extend their leadership. In addition to absorbing the greatest share of data scientists today, technology and financial services companies are planning to increase their investment in AI by the greatest proportion in the next three years (McKinsey).

Global Institute). Technology and financial service companies are prioritising AI, committing resources and building network effects around people and data to establish and extend leadership in the field.

Conversely, select sectors – including retail and consulting – are lagging, both in their ability to attract AI talent today and in their investment for the future. While many sectors, including retail and consulting, offer numerous prediction and optimisation problems well suited to AI, and large data sets to train AI algorithms effectively, the emerging gap between winners and losers in the war for AI talent is likely to widen.

The 'brain drain' from academia is real

The perceived 'brain drain' from academia to industry is real – and will have mixed implications. While alternative surveys suggest that up to 15% of data scientists currently work in academia (Kaggle), many are leaving for roles in global technology companies. A three- to five-fold increase in salary, vast data sets for analysis and access to greater hardware resources attract many. Between 2006 and 2014, the proportion of AI research publications including an author with corporate affiliation increased from approximately 2% to nearly 40% (The Economist). Talent has continued to migrate to industry. In the UK, in the last 18 months several leading AI researchers have moved to industry to accept senior roles at Uber, Amazon and Google.

In industry, AI experts are freed from the burden of securing research grants, may innovate faster, and can catalyse AI's immediate impact on the world. However, their migration

has drawbacks – including fewer teachers to train the next generation of practitioners, a concentration of expertise and experience in a small number of companies, and reduced sharing of ideas. National talent working for the public good is becoming overseas resource for private gain – with international implications. The field of AI itself arose from academic experimentation. If we lose the next generation of academics, "in the end, society will suffer" (Maja Fritsch, Professor of Affective and Behavioural Computing, Imperial College London).

High job satisfaction is intensifying the war for talent

Competition for AI talent is fierce, not simply because supply is limited. Three quarters of AI developers are content with their current roles, rating their job satisfaction 6 out of 10 or better (Fig. 69).

To optimise hiring and retention, companies should align roles to AI professionals' primary motivators. To developers, opportunities for learning and professional development, the office environment in which they will be working, and the technologies (languages and frameworks) they will be using are more important than money (Fig. 70).

Fig 69. Three quarters of AI developers are satisfied (6 out of 10 or better) with their current roles





Fig 70. Learning, office environment and the technologies they will use are AI developers' primary motivators



Source: Kaggle

Large companies seeking to attract AI talent should: take advantage of their ability to pay high salaries and offer job security; highlight the large data sets they have for analysis and the learning opportunities these will provide; emphasise the impact AI developers will have given the companies' large customer bases; and offer their AI professionals extensive hardware and software resources. Large companies should seek to mitigate likely concerns regarding agility, autonomy and freedom to publish.


Startups and scale-ups cannot, and need not, compete with the pay offered by large companies. Startups should market to candidates the intellectual and technical challenges they can provide and associated learning opportunities; an engaging office environment; impressive job titles; a greater opportunity to impact product; increased autonomy; faster cycles of innovation; and greater freedom to publish. Startups should address probable concerns regarding pay by highlighting the large, long-term financial rewards they can offer through equity awards.

Start-ups and scale-ups cannot, and need not, compete with the pay offered by large companies.

New practitioners are following sub-optimal paths to employment


The pathways into AI employment – company websites and technology job boards – prioritised by those entering the field are among the least effective (Fig. 71). People successfully employed in AI highlight that engagement with recruiters, friends, family and colleagues is the most fruitful route into the industry.

Fig 71. The most effective route into AI work is engagement with recruiters



Source: Kaggle

The pathways into AI employment
prioritised by those entering the field –
company websites and technology job
boards – are among the least effective.



Chapter 7



Europe's AI startups

The landscape for entrepreneurs is changing. Europe's 1,600 AI startups are maturing, bringing creative destruction to new industries, and navigating new opportunities and challenges. While the UK is the powerhouse of European AI, Germany and France may extend their influence.

Summary

- Europe is home to 1,600 early stage AI software companies. AI entrepreneurship is becoming mainstream. In 2013, one in 50 new startups embraced AI. Today, one in 12 put it at the heart of their value proposition.
- The European startup ecosystem is maturing. One in six European AI companies is a growth-stage company with over \$8m of funding. Expect acquisitions to recycle capital and talent; startups competing with 'scale-ups' as well as incumbents; and increasing competition for talent.
- The UK is the powerhouse of European AI with nearly 500 AI startups – a third of Europe's total and twice as many as any other country. We provide a map of the UK's AI startups and feature 14 leading companies.
- Germany and France are thriving European AI hubs. High-quality talent, increasing investment and a growing roster of breakout AI companies are creating feedback loops of growth and investment.
- Spain's contribution to European AI exceeds its size. Immigration, which correlates with entrepreneurship, has deepened the Country's talent pool.
- The European AI landscape is in flux. While the UK remains the powerhouse of European AI, its share of European AI startups, by volume, has slightly reduced. Brexit could accelerate this. France, Germany and other countries may extend their influence in the decade ahead, spreading the benefits of entrepreneurship more evenly across Europe.
- Italy, Sweden and Germany 'punch above their weight' in core AI technology, while there is support for Nordic countries' reputation for deep tech expertise.
- Nine in ten AI startups address a business function or sector ('vertical'). Just one in ten provides a 'horizontal' AI technology.
- A quarter of new AI startups are consumer companies, as entrepreneurs address or circumvent the 'cold start' data challenge. Many focus on finance or health & wellbeing.
- Healthcare, financial services, retail and media & entertainment are well served by AI startups. In sectors including manufacturing and agriculture, entrepreneurial activity is modest relative to market opportunities.
- Health & wellbeing is a focal point for AI entrepreneurship; more startups focus on the sector than any other. In the coming decade, developers will have a greater impact on the future of healthcare than doctors. Activity is thriving given profound new opportunities for process automation and a tipping point in stakeholders' openness to innovation.
- The UK is the heartland of European healthcare AI, with a third of the Continent's startups. UK entrepreneurs benefit from healthcare scale-ups stimulating talent and increasing openness to innovation within the NHS.
- Marketing and customer service departments enjoy a rich ecosystem of suppliers. A quarter of AI startups serving a business function focus on marketing teams.
- An influx of AI startups supporting operations teams is driving increasing process automation.
- AI companies raise larger amounts of capital, due to technology fundamentals and extensive capital supply.
- Core technology providers attract a disproportionate share of funding. While comprising a tenth of AI startups, they attract a fifth of venture capital.
- AI entrepreneurs' key challenges are the availability of talent, access to training data and the difficulty of creating production-ready technology.

Recommendations

Executives

- A growing proportion of startups and scale-ups are using AI to deliver new capabilities. Assess the extent to which your suppliers are taking advantage of AI.
- Use our market map to explore the rich ecosystem of early stage UK companies putting AI at the heart of their value proposition. Most are B2B vendors and some will offer market-leading solutions to challenges in your organisation.
- Take advantage of the influx of new suppliers serving operations teams to reassess the potential for process automation in your organisation.
- Early stage AI companies value the training data, and testimonials, your organisation can provide. Suppliers may be willing to adapt their pricing, or solution, to your requirements in return for your data and public endorsement.

Entrepreneurs

- In a crowded market, prioritise customer acquisition over short-term revenue to take advantage of data network effects that enable long-term differentiation.
- Identify potential competitors and partners in the UK using our market map.
- If beginning a venture, explore functions and sectors where activity is limited relative to market opportunity, including agriculture and manufacturing.
- Europe's AI ecosystem is maturing. If yours is a later stage company, leverage product maturity, customer references and capital to secure competitive advantage. If yours is an early stage company, prioritise adaptability and speed of execution.
- To overcome challenges regarding talent, data and productising AI, read our 'AI Playbook' (immcventures.com/research) that offers best practices.
- AI companies are raising larger volumes of capital than others. Capitalise your business adequately to create and maintain competitive advantage.

Investors

- With select sectors and functions over-supplied by startups, others under-served, and some witnessing an influx of new participants, identify areas of opportunity aligned with emerging dynamics and the themes on which you focus.
- As AI startups mature, evaluate opportunities to support portfolio companies with emerging challenges including international expansion and acquisitions.
- With investments into AI companies larger than average, valuations can be elevated. Consider whether or not you are willing to 'overpay' to access opportunities.

Policy-makers

- Competition for talent and capital is increasingly pan-European. Support your country's early stage companies by removing impediments to the flow of skilled talent and international capital.
- Expand public sector organisations' openness to innovation, and simplify procurement processes, to catalyse opportunities for early stage companies and deliver improved public services.

Europe is home to 1,600 early stage AI software companies

With every paradigm shift in technology, innovative early stage companies emerge to improve and then reimagine business processes and consumer applications.


Over time, the distinction between 'AI companies' and other software providers will blur and then disappear, as AI becomes pervasive. Today, however, it is possible to highlight a sub-set of early stage software companies that have AI at the heart of their value proposition.

We individually reviewed the activities, focus and funding of 2,830 purported AI startups in the 13 EU countries most active in AI – Austria, Denmark, Finland, France, Germany, Ireland, Italy, the Netherlands, Norway, Portugal, Spain, Sweden and the United Kingdom. Together, these countries also comprise nearly 90% of EU GDP. In approximately 60% of the cases – 1,590 companies – there was evidence of AI material to a company's value proposition.

AI entrepreneurship is becoming mainstream

In 2013, just one in 50 new startups embraced AI. Today, one in twelve put AI at the heart of their value proposition (Fig. 72). In 2019, entrepreneurs are disrupting incumbents by leading the paradigm shift to AI.

Fig 72. One in twelve new European startups is an AI company



Source: M&A Ventures (2018 data to October)

AI-led startups have proliferated since 2016, as technological enablers for AI meet triggers for entrepreneurship. Maturing AI enablers included: enhanced algorithms offering improved results; specialised hardware that accelerated AI system training; and greater availability of training data.

Against this background, more entrepreneurs are taking advantage of AI as: cloud-based AI infrastructure and open source AI frameworks reduce initiation and scaling costs; startups successfully access pools of AI talent at leading universities; venture capital funding for European AI startups has increased as providers of capital recognise opportunity for returns; and successful AI exits (Blue Vision Labs, Deep Mind, MagicPony, SwiftKey) and scale-ups (including Ada Health, Babylon Health, Benevolent AI, Darktrace, Graphcore, Kreditech and Meero) highlight demand and recycle capital and leadership experience within the European ecosystem.

Within ten years, most companies will use AI in select business processes, either directly or via their suppliers. Widespread adoption of AI among today's entrepreneurs is a leading indicator of a near-term future in which AI is pervasive.

For incumbents, the growth of AI entrepreneurship is a double-edged sword. AI startups are valuable suppliers – an ‘on-ramp’ to AI – for companies that embrace them, while disrupting those that do not. Select early stage companies will be acquired by today's incumbents or become the incumbents of tomorrow.

The European AI ecosystem is maturing

While AI entrepreneurship is nascent (six in ten AI startups in Europe are at the earliest stages of their journey, with Angel or Seed-stage funding), it is maturing. One in six European AI companies has passed through Angel, Seed and Early Stage phases to a ‘Growth’ phase catalysed with over \$8m in venture funding (Fig. 73).

Countries with a large number of AI companies (the UK, France and Germany) typically have more mature ecosystems (Fig. 74). In the UK, France and Germany, one in five AI startups are later, ‘Growth’-stage companies; in Sweden, just one in ten. Spain is an exception. While there are almost as many AI companies in Spain as in Germany, just one in ten is mature.

As the ecosystem matures, we expect:


- an increasing number of exits, as incumbents acquire disruptive, early stage companies gaining critical mass;
- a positive, ‘flywheel’ effect as lucrative exits recycle capital and talent within the ecosystem;
- selective, high-profile failures among companies that have raised significant sums of capital;
- startups competing with ‘scale-ups’ as well as incumbents;
- increasing competition for technical talent and executive leadership, as ‘scale-ups’ offer attractive salaries and impact as well as innovation;
- the European AI sector to better compete with larger US vendors as, to an increasing extent, multinational companies procure vendors internationally.

Fig 73. Six in ten startups are at the Angel or Seed stages



Source: MMC Ventures, Bearhunt, Crunchbase, Traction

Fig 74. Larger ecosystems are typically more mature; Spain is an exception



Source: MMC Ventures, Bearhunt, Crunchbase, Traction

The UK is the powerhouse of European AI

Fig 75. With twice as many AI startups as any other country, the UK is the powerhouse of European AI entrepreneurship



Source: MMC Ventures, Beahunt, Crunchbase, Tracm

With nearly 500 AI startups – a third of the European total and twice as many as the next most active country – the UK is the heartland of European AI (Fig. 75). With the largest internet economy in the G20, extensive academic talent including a quarter of the world's top 25 universities, a growing number of AI exits (DeepMind, SwiftKey, MagicPony) recycling capital and talent, supportive Government policy in relation to AI, and a global financial services hub, the UK has significant assets.

The market map, overlaid, places the UK's 500 startups according to:

- **Purpose:** Does the company focus on a business function (for example, marketing or human resources), a sector (healthcare, education) or core AI technology with cross-domain application?
- **Customer:** Does the business predominantly sell to other businesses (B2B) or to consumers (B2C)?
- **Funding:** How much funding has the company disclosed to date? We categorise companies as: Angel or Seed stage (under \$500,000 to \$2m); or Early or Growth stage (over \$2m to c. \$200m).

With approximately 200 AI startups each, Germany and France are thriving AI hubs in Europe. High quality talent, increasing volumes of capital and an expanding roster of successful AI companies are creating feedback loops of growth and investment.

Spain is an outlier whose contribution to European AI exceeds its size. Despite a population half the size of Germany, Spain houses almost as many AI startups. Extensive immigration may have deepened the Country's already broad pool of talent. Spain has the second highest rate of immigration in the EU, and entrepreneurial activity is higher among immigrants than native citizens (Global Entrepreneurship Monitor).

KEY

- B2B ■ B2C
■ New company (since 2017)

Funding category

- Angel/Seed: <\$2m
 Early stage/Growth: >\$2m

UK AI Landscape (Early stage companies)

Core Technologies

AI INFRASTRUCTURE

EARLY STAGE / GROWTH
 Fetch AI ■ Oran Labs ■ Graphcore ■ Hadian ■ Inviso ■
 Serion ■

ANGEL / SEED
 Brydyl ■ Hivemind ■ Insilabs ■ Monograph ■
 Skye Technologies ■

AUTONOMOUS SYSTEMS

EARLY STAGE / GROWTH
 FiveAI ■ Focal Point Positioning ■ Latent Logic ■ Oxbeston ■ Provionix ■
 Sky Futures ■

ANGEL / SEED
 Academy of Robotics ■ Accelerated Dynamics ■ Banci Vehicles ■ Evolve Dynamics ■
 Headlight AI ■ Humanizing Autonomy ■ Intelligent Robots ■
 Machines With Vision ■ Prodrive ■ React AI ■ Third Space Auto ■ Wayve ■

COMPUTER VISION

EARLY STAGE / GROWTH
 Blue Vision Labs ■ Microblink ■ SLAMcore ■
 Spectral Edge ■ Visio ■

ANGEL / SEED
 Aithena Technologies ■ Emteq ■ Facefish ■
 Machine Medicine ■ Synthesia ■ TouchByte ■
 Visio Igenix ■ Vize.ai ■

LANGUAGE & SOUND

EARLY STAGE / GROWTH
 Audio analytic ■
 Procto.ai ■ Speechmatics ■
 Telicor ■

ANGEL / SEED
 Capito Systems ■ Kami.ai ■
 Rely.ai ■

OTHER

EARLY STAGE / GROWTH
 Cambridge Quantum Computing ■
 Ditto AI ■ iFlow ■
 Mind Foundry ■

ANGEL / SEED
 Both co Chatbot Agency ■

Sector

AGRICULTURE

EARLY STAGE / GROWTH
 KismetHub ■ Doughty Technologies ■
 Global Surface Intelligence ■
 Hummingbird Technologies ■
 Optimal Labs ■ Xihai ■

ANGEL / SEED
 Converse ■ Henspace ■ Observe ■

EDUCATION

EARLY STAGE / GROWTH
 Bridge17 ■ Century ■ Lingvist ■
 Lingvist ■

ANGEL / SEED
 Gernan ■ Reizai ■ Qeademy ■ Synap ■

FINANCIAL SERVICES

EARLY STAGE / GROWTH
 Arkana ■ BML Technology ■ Clio AI ■ Cytura ■ Digital Contact ■ Esense Analytics ■ Floodfish ■ ForwardLane ■ FraternityScore ■ Kivardach ■
 Mosaic Smart Data ■ Multiply ■ Plun ■ PriceHubbly ■ Proccurus ■ Quantemplate ■ Quantina ■ Tractable ■ TradeIQ ■ Virtuous ■

ANGEL / SEED
 Alaska ■ Acuity Trading ■ AdviceGames ■ AlgoDynamics ■ Alpha+ ■ ArrayStream ■ Artificial Labs ■ Banked ■ Blue Lion Research ■ Broly ■ Burd ■
 Chip ■ CLUE ■ Dealr ■ Digital Clipboard ■ Ducit.ai ■ Enlago ■ Financial Network Analytics ■ G-Research ■ Kneer EA ■ Metakood ■ Opportunity ■
 Providentia ■ Rightthink ■ Sns Analytics ■ Spol ■ Tonalogy Tech ■ TradeRise ■ Win FX Financial Innovation ■ Yedoo ■ Zora ■

FOOD & BEVERAGE

EARLY STAGE / GROWTH
 Gousto ■ Henschen ■
 Nuritas ■ Winnow ■

ANGEL / SEED
 Dinadite Limited ■
 Intelligent ■ Vidualia ■

HEALTH & WELLBEING

EARLY STAGE / GROWTH
 Amiko ■ Babylon Health ■ BenevolentAI ■ BioBeats ■ BrainWaveBank ■ Cambridge Medical Robotics ■ Cydar Medical ■ Deontics ■ Drayson Technologies ■
 Excelsia ■ GFN ■ HealthUnlocked ■ Heals ■ Kheiron ■ Kraydon ■ LabGenius ■ Libell ■ Medsped ■ Snap4D ■ Synthace ■ Thrua ■
 Visulytic ■ Via ■ Your MD ■

ANGEL / SEED
 Algenpuls ■ Antivene ■ Auro ■ Bioline ■ Cambridge Bio-Augmentation Systems ■ Cambridge Humane ■ Causaly ■ Chronomics ■ eCURE ■
 FitWeb ■ InvenioLabs ■ InsideDNA ■ Kalido ■ Kinoku ■ ONE Health ■ Opasium ■ Peptone ■ PetaGene ■ Resurgo Genetics ■ Sentinel ■
 Sine Diagnostics ■ ThinkSono ■ Transformative ■

INFRASTRUCTURE & TRANSPORT

EARLY STAGE / GROWTH
 Energi Mine ■ Limejump ■ nPlan ■
 Serdat ■ Vero by Green Running ■

ANGEL / SEED
 Calluna ■ CrowdConnected ■ Dispense.io ■
 Green Running ■ Grid Edge ■ NumberEight ■
 OpenCapacity ■ Osven Telematics ■ Rocco ■
 TravelAI ■ Wupper ■

LAW

EARLY STAGE / GROWTH
 ClauseMatch ■
 Eigen Technologies ■
 Lumina ■ ThoughtRiver ■

MANUFACTURING

EARLY STAGE / GROWTH
 ClouMAC ■ Neuron soundways ■
 QIO Technologies ■ Senseye ■

MEDIA & ENTERTAINMENT

EARLY STAGE / GROWTH
 Instrumental ■ Jukebox ■ Myrtle Software ■ Mogens ■ Pappa ■ Spirit AI ■

PROPERTY & REAL ESTATE

EARLY STAGE / GROWTH
 Papi ■

RETAIL

EARLY STAGE / GROWTH
 Cocoon ■ Cortexica ■ Dressipi ■ EDITED ■ Enventech ■ Mallo ■ Metal ■
 Snap Tech ■ Spoon Guru ■ Thread ■

Function




Germany and France are extending their influence

The dynamics of AI entrepreneurship in Europe are in flux.

While the UK remains the powerhouse of European AI, and will house more AI startups than other European countries for years to come, its share of European AI startups, by volume, has slightly reduced (Fig. 76). Brexit could accelerate this dynamic. AI developers are skilled, few in number and may select opportunities from the many offers they receive. More broadly, one in five London technology workers is an EU national from overseas (London Tech Advocates). If free movement of workers between the EU and UK ends, visas are unforthcoming, or rhetoric is unwelcoming, the UK's access to talent could reduce. France, Germany and other countries may extend their influence in the decade ahead, spreading the benefits of entrepreneurship more evenly across Europe.

The dynamics of AI entrepreneurship are in flux.
While the UK remains the powerhouse of European AI, other countries may extend their influence in the decade ahead.

Fig 76. France and Germany are increasing their share of European AI entrepreneurship




Italy, Sweden and Germany 'punch above their weight' in core technology

While two thirds of Europe's core technology AI startups are located in the UK, Germany, Spain and France, adjusting for countries' size – their number of AI startups – reveals a different dynamic:

Relative to their size, Italy, Sweden and Germany are core technology hubs; in each, approximately one in five AI startups is a core technology provider compared with the European average of one in eight (Fig. 7.7). There is also support for Nordic countries' reputation for deep tech expertise: in Finland, Denmark and Norway one in seven AI startups is a vendor of core AI technology.

While countries with large AI ecosystems, such as the UK, benefit from a large number of leading universities, broad pools of talent and extensive investment, smaller hubs 'punch above their weight' for varying reasons. In addition to exceptional talent, their ecosystems benefit from: leading research and engineering centres (Germany); effective core technology incubators (Finland); the AI laboratories of internet giants (France); and the 'halo' effect of multiple successful scale-ups in other fields (Sweden). Flows of venture capital into smaller core technology hubs are also increasing, creating a virtuous circle of investment and success.

Fig 7.7. While many core technology startups are in the UK or Germany, relative to their size, Italy, Sweden and Germany are core technology hubs




Source: IAVIC, Venturis, Benchmark, Crunchbase, PwC

Nine in ten AI startups are B2B

Nine in ten of Europe's 1,600 AI startups are business-to-business (B2B) vendors, developing and selling solutions to other companies (Fig. 78). Just one in ten sells directly to consumers (B2C).

Historically, B2C AI has been inhibited by the 'cold start' data challenge. Training AI algorithms typically requires large volumes of data. While B2B companies can analyse the extensive data sets of the businesses they serve, customer-facing companies usually begin without large volumes of consumer data to analyse (in the absence of public or permissioned data, such as Facebook profile information). B2C companies typically deploy AI later, as their user bases and data sets grow.

Fig. 78. Nine in ten AI startups are B2B




Source: M&C Ventures, Statista, Crunchbase, Techn

B2C AI is on the rise as the 'cold start' thaws

While most existing AI companies are B2B, a growing proportion of new AI startups – in 2018, a quarter – are B2C (Fig. 79). B2C AI startups are mitigating or circumventing the cold start challenge.

From their inception, a greater proportion of new B2C companies are planning effective data acquisition strategies for AI. By integrating with existing customer data (such as

Fig. 79. B2C AI is on the rise – a quarter of new AI companies are B2C




Source: M&C Ventures, Statista, Crunchbase, Techn

financial transaction information), capturing data earlier in their customers' journeys, or developing partnerships with data providers and other companies, companies are mitigating the cold start challenge to gain value from AI earlier in their lives. While incumbent consumer companies struggle with sprawling, siloed data estates, AI startups are turning a limitation to an advantage by creating a data collection and processing pipeline optimised for AI.

Entrepreneurs are also circumventing the challenge by imaginatively applying AI techniques to a wider range of consumer processes. Without extensive third-party data sets, early stage consumer companies can present new forms of engagement (such as human-computer interaction via chatbots) and offer new services and experiences (by using AI to optimise their supply chain).


The rise of B2C AI also reflects a shift in entrepreneurship to B2C-leaning sectors. There is a higher proportion of B2C AI companies in which data is more readily available: media & entertainment (47% B2C), finance (26%), and health & wellbeing (27%) (Fig. 80). In the last 24 months, the sectors attracting the highest proportion of new AI startups have been: finance (23% of new startups), health & wellbeing (17%), and media & entertainment (10%) (Fig. 81). As entrepreneurs tackle B2C-leaning sectors, B2C AI is on the rise.

Fig 80. There is a higher proportion of B2C AI companies in sectors where data is readily available



Source: MMC Ventures, Beahunt, Crunchbase, Tracx

Fig 81. Half of new AI startups target the finance, health or media sectors



Source: MMC Ventures, Beahunt, Crunchbase, Tracx

AI entrepreneurship remains vertically-focused

Nine in ten AI startups address a need in a specific 'vertical' (business function or sector) (Fig. 82), just one in ten is developing a core, 'horizontal', AI technology (a sector-agnostic capability or platform). This mix has remained consistent over time (Fig. 83).

The proportion of core technology providers will remain modest. Google, Amazon, IBM and Microsoft (GAIM) offer an extensive, and expanding, suite of core AI technologies, primarily in the fields of computer vision and language. Their solutions – ranging from audio transcription, language translation and sentiment analysis to object recognition and facial analysis – are capable and leave limited room for any but the most specialised direct competitors. Further, developing core technology requires world-class technical expertise (frequently stemming from academic research) which is limited in supply.

GAIM solutions, however, are generic and sector-agnostic. AI startups are addressing the myriad sector- and function-specific opportunities which GAIM vendors lack the strategic desire, domain expertise and data advantage to address.




Fig 82. One in ten AI startups focuses on core tech

Source: MWC Ventures, Beaufort, Conchise, Ticon

Fig 83. The proportion of 'horizontal' core technology providers has remained consistent over time



Source: MWC Ventures, Beaufort, Conchise, Ticon

The healthcare and financial services sectors are well served by AI startups

The health & wellbeing, finance, retail and media & entertainment sectors are well served by AI startups (Fig. 84). Activity in these sectors is high, in part, because they are well positioned to benefit from AI technology while offering attractive commercial characteristics for entrepreneurs. Active sectors offer:

- Large market opportunities with domain-specific challenges unaddressed by the generic AI offerings of platform vendors Google, Amazon, IBM and Microsoft (GAIM).
- numerous prediction and optimisation challenges well suited to the application of AI;
- large data sets for training and deployment, although access to data in healthcare can be challenging;
- a path to better-than-human performance, through AI, that is technically achievable;
- opportunity for significant, demonstrable value creation, such as improved trading performance (financial services) or improved purchase conversion (retail);
- alternatives to automation that are impractical (healthcare) or expensive (finance).

Fig B4. More AI startups – one in five – serve the health & wellbeing sector than any other



Source: MMC Ventures, Beahut, Crunchbase, PwC

In select areas, activity is modest relative to market opportunities. In manufacturing, few startups address a substantial need. Manufacturers could reduce material costs with improved analysis of product quality. Buffering (the storage of raw materials to compensate for unforeseen production inefficiencies) could be reduced by up to 30% with more predictable production. The requirement for significant domain expertise serves as an inhibitor to younger entrepreneurs in this area.

In other sectors, such as education, activity is inhibited by technology fit (stakeholders spend a lower proportion of time collating and processing data – 23% in education versus 50% in finance) and commercial considerations (challenging buyer dynamics).

Healthcare is a focal point for AI entrepreneurship

More AI startups – one in five – serve the health & wellbeing sector than any other (Fig. B4). In the coming decade, developers will have a greater impact on the future of healthcare than doctors. Healthcare is a focal point for AI entrepreneurship as:

- AI offers profound new opportunities for process automation and cost reduction in healthcare, as AI technologies (computer vision, natural language processing and improved pattern matching) enable formerly human processes to be undertaken in software at scale and low cost. AI can improve most stages of an individual's healthcare journey (including diagnosis, treatment and monitoring) and associated workflows (triage, drug discovery and fulfillment);
- challenges to healthcare systems reach a 'tipping point': Ageing populations and new medical treatments are increasing costs. In many European countries, since 1970 healthcare costs as a percentage of GDP have doubled to approximately ten per cent (OECD). Further, as austerity pressures governments' spending, consumer expectations continue to rise;
- increasingly, healthcare system stakeholders are willing to embrace innovation and early stage companies (in October 2018, the UK Health Minister published a vision for the future of UK healthcare with modern technologies at its core);
- the already vast market opportunity in healthcare expands with the rise of wellbeing-related applications (fitness, meditation, talking therapies and preventative testing); and
- a cohort of bold entrepreneurs, many who combine medical expertise with commercial acumen, seek to effect structural change at scale.

The UK is the heartland of European healthcare AI

With one in three of the Continent's startups, the UK is the heartland of European healthcare AI. In addition to having more AI startups, overall, than any other European country, and larger quantities of venture capital investment, UK healthcare entrepreneurs benefit from:

- Many of the world's top-rated universities for medicine, and teaching hospitals, that create a large pool of expert practitioners and opportunities for collaboration between researchers, startups and care providers;
- the 'flywheel' effect of a critical mass of healthcare scale-ups. Companies including Babylon Health, Benevolent AI, DeepMind Technologies and Sophia Genetics are stimulating, attracting and recycling talent, capital and commercial engagement in the UK ecosystem;
- Increasing openness to innovation in the NHS. 'The tech revolution is coming to the NHS' (UK Health and Social Care Secretary). While engaging with the NHS remains challenging given its scale, fragmentation and procurement procedures, early stage companies are benefiting from more accessible deployment opportunities as the Government seeks to 'transform the NHS into an ecosystem of enterprise and innovation that allows technology to flourish and evolve' and to establish 'open standards' (Department for Health and Social Care).

- a Government commitment to increase the budget of NHS England and above inflation by an average of 3.4% each year until 2023/24, and policies to catalyse healthcare AI including a £50m investment in five new AI medical technology centres in 2019.

There remain inhibitors and sources of uncertainty for healthcare innovation in the UK – including disparate data standards and conflicting IT systems within the NHS, unclear data permissioning protocols, budget pressures in areas including social care, and Brexit.

Marketing and customer service teams enjoy a rich ecosystem of suppliers

Marketers are well served by Europe's AI entrepreneurs. Among AI companies serving a business function, more – a quarter – focus on marketing departments than any other. Customer service and IT departments also receive significant attention (one in six startups, respectively) (Fig. 85).

While the UK contributes half of Europe's AI marketing startups (Fig. 86), France is Europe's hub for AI customer service with a fifth of the Continent's startups (Fig. 87).

Fig 85. More startups – one in four – serve the marketing function than any other






Fig 86. The UK contributes four in ten of Europe's AI marketing startups (% of European AI marketing startups)



Source: MMC Ventures, Bainhurst, Churchill, Tracen

Fig 87. France is Europe's hub for AI Customer Service (% of European AI Customer Service startups)




Source: MMC Ventures, Bainhurst, Churchill, Tracen

Modern marketing represents a sweet-spot for AI. Consumers have billions of touch points with websites and apps, providing a rich stream of complex data that is difficult to analyse using traditional, rules-based software but well suited to AI-powered analytics. In addition, natural language AI enables supplementary data, such as social media, to be analysed at scale for the first time. Most stages of the marketing and advertising value chain are ripe for optimisation and automation, including: consumer segmentation; consumer targeting; programmatic advertising; consumer purchase discovery; and consumer sentiment analysis. Competition and commoditisation are primary challenges for early stage AI marketing and advertising companies.

Customer Service departments are well served following a recent wave of new, AI-powered vendors. Among those addressing a business function, one in five AI startups founded since 2017 sell customer service solutions (Fig. 88). Entrepreneurs are taking advantage of advances in natural language processing AI to offer new augmented or automated customer service capabilities including: social listening (identifying and responding to customers automatically); intelligent classification and routing of contact centre enquiries; drafting or full automation of contact centre responses; chatbots (for customer engagement); and automated customer care analytics.

Modern marketing represents
a sweet-spot for AI. Consumers
have billions of touch points with
websites and apps, providing a
rich stream of complex data well
suited to AI-powered analytics.

Fig. 88. Among new startups addressing a business function, one in seven serve operations teams



Source: BMC Software, Kantar, Crunchbase, Ycombi

An influx of AI startups is driving process automation

While currently underserved, the operations function is benefitting from an influx of new, AI-led startups in the last 24 months. Among those addressing a business function, one in seven AI startups founded since 2017 serve operations teams (Fig. 88).

Traditional data mining techniques are less effective for process control given systems' varying media and data formats, concurrency, loops and decision-making (Chabunok). Advances in AI computer vision, natural language processing, understanding and reasoning are expanding the breadth of materials accessible to digital automation, offering greater understanding of their content, and enabling more intelligent responses.

AI is profoundly expanding the envelope of automation – the breadth and value of processes susceptible to digital mechanization. Improved capabilities include: recommending the next best action in a workflow; better automation of document processing; and more expansive robotic process automation (RPA). In the short term we expect the number of vendors serving the Operations function to increase further. In the medium term, commoditization and competition will become challenges. Vendors focusing on a particular industry may develop the domain expertise, deep workflow integrations, data network effects and referencability to develop lasting competitive advantage.

AI companies raise larger investment rounds


Since 2015, when securing investment AI companies have raised larger volumes of capital than traditional software companies (Fig. 89). A difference exists across all stages of maturity, from Seed stage through Series A, B and C funding (Fig. 90).

Early stage AI companies are attracting larger funding rounds, due to sector fundamentals and dynamics in the supply and demand of capital.

AI companies' capital requirements can justify greater investment, given the longer cycles required to achieve, develop a minimum viable product, the high cost of AI talent, and the larger teams required for complex deployments.


Beyond fundamentals, capital influsions are being inflated by extensive supply of capital and limited demand. Many venture capitalists wish to invest in AI but there are relatively few AI companies in which to invest. Globally, venture capital investment in early stage AI companies has increased 15-fold in five years, while the number of investable prospects remains limited. As the number of AI-led startups has increased (today, one in 12 new startups in Europe is an AI-led startup) differences in round sizes are reducing.

Fig 89. Since 2015, AI companies have raised larger investment rounds



Source: Crunchbase, IAMC, Venture

Fig 90. AI companies are raising larger rounds at all stages of maturity




Source: Crunchbase, IAMC, Venture

AI companies attract premium valuations

AI companies are securing higher valuations, as well as securing larger capital infusions. Distributing companies founded since 2016 along a valuation curve reveals that a smaller proportion of AI companies are valued at lower amounts, and a greater proportion are valued at higher amounts, than equivalent non-AI startups. This is the case across most stages of maturity (Fig. 91) and within the early phases a company's life (Fig. 92).


Pragmatically, entrepreneurs raising large volumes of capital seek higher valuations to avoid unpalatable ownership dilution. Investors may also be willing to value highly AI companies that have attracted scarce AI talent, developed advanced and defensible technology, or have a data advantage delivering data network effects. Beyond industry fundamentals, an imbalance in demand for capital and its supply is inflating valuations. AI companies' valuations benefit from investors competing to deploy capital into a limited number of AI prospects. With AI entrepreneurship becoming mainstream (page 99), this tailwind will reduce.

Fig 91. A greater proportion of AI startups are highly valued



Source: Distribution.co, M&A Ventures

Fig 92. A greater proportion of AI startups are highly valued (detail – lower values)



Source: Distribution.co, M&A Ventures

Core technology providers attract a disproportionate share of funding

Core technology providers – ‘deep tech’ companies developing ‘horizontal, sector-agnostic capabilities instead of ‘vertical’ solutions focused on individual sectors or business functions – attract a disproportionate share of venture capital (Fig. 93). While core technology companies comprise a tenth of AI startups, they attract a fifth of venture capital investment.


Core technology companies, from developers of autonomous systems to computer vision and language companies, exhibit more fully the capital dynamics latent in AI:

- developing core technology demands expensive, world-class talent;
- the time required to develop minimum viable products can be longer in the technically demanding field of core AI, increasing cash burn; and
- a greater proportion of core technology companies pursue atypical revenue models, such as licensing agreements, in place of traditional licenses or software-as-a-service subscription agreements, elongating time to revenue.

Core technology entrepreneurs should adequately capitalise their businesses for longer, deeper periods of expenditure, while their investors develop syndicates with deep pockets. Doing so can enable core technology companies to realise their potential: capturing vast market opportunities with differentiated, defensible technology.

While core technology companies comprise a tenth of AI start-ups, they attract a fifth of venture capital investment.

Fig 93. Core technology companies attract a disproportionate share of funding



Talent, data and production are AI startups' key challenges

Competition for AI talent, the limited availability of training data, and the difficulty of creating production-ready technology are consistently entrepreneurs' key challenges when developing AI.

1. Recruiting AI talent is challenging

Startups compete with multiple categories of competitors – including large technology companies (Google, Amazon, IBM, Microsoft, Facebook), banks, professional service firms, and other early stage companies – for data scientists, AI experts and AI engineers. Recruiting staff that have a balance between theoretical expertise and commercial experience, and experience running an AI team, are additional difficulties.



"Access to talent, and its competitiveness, is the biggest challenge."
David Benigson, Signal



"London is a good place to be, when looking for AI talent."
Dmitry Alexseyev, DigitalGenius



"London has one of the best pools of AI talent in the world – which is the main reason why we are here."

Fabio Kuhn, Vortexa

To identify and attract talent, AI-led companies are building deep relationships with academic institutions, being active members of research communities, publishing papers, and collaborating with universities.



"We try to engage with developers well before they're looking for a job, and let them do what they love."

David Benigson, Signal

2. Access to training data is critical

Access to initial data sets for training poses a challenge:



"It's a classic chicken and egg problem. Early customers, and thus data, are hard to get when you don't have any existing reference clients."

Tim Sadler, Tessian

Companies are mitigating the difficulty by developing powerful use cases for access to client data and by implementing a data acquisition strategy from early in their lives.



"We started collecting data very early in our journey."

Tim Bould, Gousto

For many early stage AI companies, compromising on early pricing to secure access to valuable customer data is proving effective.

3. Developing production-ready AI is difficult

Entrepreneurs recommend moving from 'lab to live' as soon as possible, testing development systems on low-risk real world data. Cross-functional collaboration is also key.



"Taking what works well in a lab and getting it to work in a diverse and sick population is a big challenge."


Chris McCann, Current Health



"The real world is full of black swans and exceptions. We've learned to overcome them by getting great at cross functional collaboration, building integration with the tech team, and constant monitoring of risk."

Time Boldt, Gousto

The UK is home to a third of Europe's AI start-ups. Overall, we feature 14 leading start-ups, spanning a range of sectors and functions, that are using AI to create new possibilities.



Featured company

Audio Analytic



Customers Businesses

Core technology Sound


We are teaching machines to “hear.” We are mapping the world of sounds - beyond speech and music - to give a wide range of devices, across a number of market sectors, the ability to understand local context.

- We give devices the ability to understand local context, making them more intelligent and helpful to users.
- Our licensees, including leading consumer technology brands, can offer unique, intelligent features in highly competitive markets.

- Our licensees can offer valuable, subscription-based services to users – for example, alerts and recordings of events at home while they are away.

We believe sound recognition is a fundamental AI technology. By 2023, sound recognition will be a “must-have” component in a wide variety of intelligent and connected devices, from smart speakers and devices to cars, mobiles and hearables. We also expect sound recognition technology to proliferate into areas beyond consumer technology.

Below (clockwise): The Audio Analytic team in Cambridge; the Free Devialet Player from pan-European telecoms operator Iliad, and the Hive Hub 360 from Hive (part of Centrica).



Featured company

Current Health



Customers: Healthcare facilities

Sector: Healthcare


We help healthcare teams stay close to the patients who need them most, at home and in hospital. Our solution is centred around an all-in-one, wireless wearable we offer, which enables remote monitoring of the human body to detect warning signs of illness and deterioration. By alerting the physician, nurse, home healthcare team or hospital to warning signs, we enable earlier, proactive intervention.

We work with hospitals, home healthcare teams and nursing facilities to:

- automate vital sign capture and data entry (in a hospital with 750 beds, five years of nursing time is spent every twelve months collecting vital sign data and inputting this into patients' medical records).

- detect deterioration earlier, to intervene sooner and save lives.
- proactively treat deterioration, to reduce (re-)hospitalisations, length of stay and the cost of healthcare.

In 1950, average life expectancy globally was 48. Today it's 70. As we age, we experience more chronic diseases, including heart failure, cancer and diabetes, which strain healthcare services. An ageing population will require a radical shift in the delivery of healthcare. We are also generating scientific and medical knowledge at a greater rate than ever before – and at a rate greater than the human mind can retain and apply. AI can scale and multiply the efforts of our doctors and nurses. With AI, our solution enables a healthcare professional to monitor hundreds of patients at once and identify the few who require help.



Featured company

Digital Genius

Customers Businesses

Function Customer Service

Our AI platform puts customer support on autopilot – by understanding conversations, automating repetitive processes and delighting customers. Our customer service automation platform is powered by deep learning, which understands customers' objectives and drives automated resolutions through APIs that connect seamlessly to a company's back-end systems. Our platform is used by KLM Royal Dutch Airlines, The Perfume Shop, Air France and other forward-looking businesses to drive conversational process automation through the use of deep learning.

Our platform can:

- automate repetitive, expensive tickets. A third of all Course Hero's customer service tickets are handled by our platform; half of these are resolved without involving an agent.
- provide shorter customer response times. With our solution, Imagine Learning's cases are answered 70% faster.
- deliver higher customer satisfaction. With our solution servicing 40% of tickets, The Perfume Shop reported a boost in customer satisfaction to 88%.

Customer service teams will evolve into customer happiness teams. Instead of fire-fighting they will focus on exception-handling and proactive outreach.



AI is transforming customer service. First, increasing automation - of conversations and processes - is inevitable given the volume of repetitive conversations in contact centres. While there will always be unique customer enquiries, increasingly sophisticated AI will increase the complexity of addressable questions. Second, customer service teams will evolve into customer happiness teams. Instead of fire-fighting they will focus on exception-handling and proactive outreach.

Teams will be measured on customer loyalty, retention, and repeat purchasing – not response time or average call duration. Finally, customer expectations regarding the quality and ease of conversations with machines will shift. Any matter that cannot be resolved in a matter of seconds will be deemed a failure in customer experience.



Business		Upgrades		[Your Tickets]	
AutoPilot		AutoPilot		AutoPilot	
ON		ON		OFF	
Days Active: 21		Days Active: 5		Days Active: 0	
Tickets Answered: 2200		Tickets Answered: 205		Tickets Answered: 0	
Cost Savings: \$19,400		Cost Savings: \$1,975		Cost Savings: 0	

Featured company

Gousto

gousto

Customers Consumers

Sector Retail

A recipe kit company, we offer consumers precise ingredients, delicious recipes and a dollop of adventure. We supply subscribers with recipe kit boxes that include ready-measured, fresh ingredients and easily followed recipes.

We offer customers:

- variety – customers can choose from 30 weekly recipes, compared with the average six to seven cooked by UK families,
- convenience – customers can order recipes in minutes, saving hours in weekly supermarket shopping,
- sustainability - we deliver ingredients in pre-portioned measures, eliminating household food waste.

Increasingly, the food industry is characterised by consumers' demands for choice and convenience. Given the complexity of meeting these expectations, AI is a vital enabler. AI powers our business, from what a customer sees on our website to how boxes are routed through our facilities. Through ever-advancing AI we will continue to differentiate our customer proposition and offer greater choice, convenience and value.

Increasingly, the food industry is characterised by consumers' demands for choice and convenience. Given the complexity of meeting these expectations, AI is a vital enabler.



Featured company

Kheiron Medical



Customers Healthcare facilities

Function Healthcare

We are a London-based AI cancer diagnostics company, with an initial focus on breast cancer screening. We believe in deep clinical rigour and robust validation – and are the first UK company to receive regulatory approval for a deep learning application in radiology.

Our first solution, Mlia, provides intelligent assessment of breast cancer screening mammograms as an independent reader. Mlia is the first and only software suited to making 'call-back' decisions as a doctor would.

The product has the potential to increase the diagnostic accuracy of population screening initiatives, reduce false positives, reduce scan-to-report times, and slash the workload of an over-stretched workforce. There is also exciting potential to bring breast cancer screening to countries where no such service currently exists.

AI in medical imaging will be a \$2bn global industry by 2023, with algorithms being integrated and deployed at an accelerating rate. Patients, radiologists and healthcare systems will benefit immensely from a plethora of diagnostic decision support tools, automation of repetitive tasks, and standardised disease screening. Kheiron is leading the way in this important field.



Featured company

Omnius

Customers: Businesses

Sector: Insurance

Our solution enables the automation of insurance claims handling, by using AI to extract structured data from varying documents and support decision-making. Our state-of-the-art AI technology allows insurers to unlock efficiencies and focus their resources on customer experience.

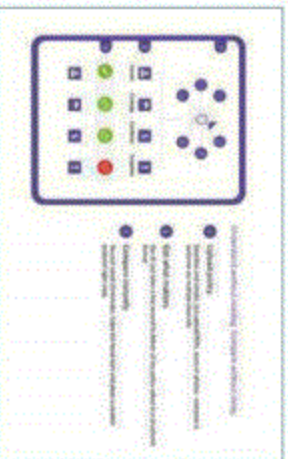
We offer:

- up to 90% automation of insurance claim handling, via cutting-edge AI which automates data extraction from varying document streams,
- deep understanding and insight into insurance, We work with eight of the 10 top global insurers,
- gains in efficiency and efficacy, by automating classification and semantic extraction of data from documents.



Insurance carriers are undergoing a once-in-a-generation technology shift, embracing new digital platforms and adopting best-in-class digital tools and analytics. The industry is moving from being process-driven to data-driven, to meet increasing customer expectations and deliver a better customer experience. AI is the game-changer that enables insurers to better understand their customers, and launch new products and services bespoke to the customer.

Insurance carriers are undergoing a once-in-a-generation technology shift, embracing new digital platforms and adopting best-in-class digital tools and analytics.



Featured company

Prowler.io



Customers Businesses

Core Technology Autonomous Systems

Our mission is to enable leaders and organisations to make better business decisions by optimising resources. Our decision engine, VUKU™, can process data in real time, adapt to uncertainty, act on sparse information and learn from experience. Our goal is to ensure that business is powered by people and empowered by AI.

VUKU unlocks benefits for businesses through better decisions

- for example:
- Lower operational costs: our platform enables 15% improvement in last mile delivery fleet utilisation.
- Greater financial returns: finance industry customers are improving the performance of their investment portfolios and delivering better returns than hedge-funds at a lower cost.


- Improved education: education industry customers are helping students to learn independently, faster. VUKU develops an automatically self-organising curriculum for students by rapidly discovering their current skill-sets and precisely tracking changes.

AI-enabled decisions will drive the world economy by 2025. New AI technologies, such as our VUKU engine, help industries optimise their business processes in ways that current technologies cannot. The expansion of AI, and the efficiencies it brings, will be as transformational for the revenues and margins of companies as computers were in the 1980s and 1990s.



Featured company

Seldon



Customers Businesses

Core Technology Tooling

Our technology accelerates the adoption of large-scale machine learning for some of the world's leading businesses. In 2017 we released Seldon Core, which has grown into one of the most popular open-source platforms for managing and deploying machine learning models. It's built on cloud-native technologies that allow models, built in any toolkit, to be run on any cloud and on-premise. Our solution is integrated into the Google Kubernetes and IBM FDL (Fabric for Deep Learning) open-source machine learning platforms.

Our solution:

- takes machine learning from research into production, providing continuous integration and delivery (CI/CD), monitoring, optimisation and compliance to deliver increased model performance and reduced risk;
- provides a uniform machine learning deployment platform across an entire organisation, to enable companies to build and deploy multi-model inference graphs across clouds and on-premise;
- streamlines model testing, deployment and optimisation workflows. Our technology enables 6x efficiency gains for faster R&D iteration and a better feedback loop between data science and DevOps teams.

Over the next five years, AI will catalyse organisations' switch from monolithic infrastructures to cloud-native, open-source stacks.

In 2011, Marc Andreessen said 'software is eating the world'. Now, AI is eating software. Over the next five years, AI will catalyse organisations' switch from monolithic infrastructures to cloud-native, open-source stacks, which leverage containers and microservices for hybrid cloud and edge deployments. New, open standards and governance frameworks will boost consumers' and regulators' confidence that model-driven decisions are accurate, explainable, and free from bias – accelerating the rate of AI adoption on an industrial scale.



Featured company

Sensat

Sensat®

Customers Businesses

Sector Construction

We use AI to solve problems that cause delays and costs

on construction sites. Our goal is to teach computers to understand the physical world in which we live. To do this, we're creating an intelligent ecosystem that simulates reality in real time, by combining real-time information and temporal data with 2D and 3D information. By better understanding our physical world, we enable computers to make intelligent decisions on our behalf – impacting the way we live and work. We were ranked by Crunchbase as Europe's top AI company in 2018.

Our visualisation platform, Mapp, allows users to measure, analyse and collaborate on their construction projects in real time.

Mapp:

- offers intelligence that automates measurement, reporting and communication to make construction sites safer and more profitable.
- improves capability, reducing the need for surveyor "boots on the ground" by up to 80%.
- increases safety, by offering asset monitoring and project management through real-time visualisation.

AI will have a profound impact on construction. Narrow intelligence, which solves individual problems well, will evolve to more general intelligence that can manage projects holistically. AI will transform safety, by removing humans from dangerous environments, and productivity, by combining on-site and off-site data to create predictive ecosystems.



Featured company

Senseon

SENSEON

Customers Businesses

Function Cybersecurity

We use AI to protect organisations from emerging cyber-attacks. We move beyond traditional rule-based applications, which are too rigid to keep pace with evolving threats, and ineffective systems that cannot differentiate between unusual behaviour and malicious activity. Our unique approach,

'AI Triangulation', understands and blends information from multiple perspectives across an organisation's entire digital estate, allowing organisations to accurately detect even the most complex, subtle cyber threats and reduce false positive alerts.

Our solution:

- reduces risk – our AI provides accurate and actionable detection that identifies cyber threats at a speed and scale that humans and legacy systems cannot.

- saves time – automated investigation frees up teams to focus on what matters: investigating genuine threats.
- saves money – our platform replaces the need for multiple single-point security solutions.

As AI capabilities have helped to improve network defence, attackers have begun to use AI for malicious purposes. The ability to evade detection or adapt attack techniques has already allowed primitive forms of machine learning-fronted attacks to breach systems and existing security tools. This is just the beginning of AI attacks – a cat and mouse game between attackers and defenders that cannot be won; instead, it is constantly evolving. Because our technology learns and adapts to change, we can keep pace with the evolution of cyber threats.



Featured company Senseye



Customers Businesses Sector Manufacturing

We offer the leading cloud-based software for industrial

predictive maintenance. Our solution helps manufacturers

avoid downtime and save money, by automatically forecasting

machine failure without the need for expert manual analysis.

Its intelligent machine learning algorithms enable it to be

used with any machine, from any manufacturer, and to absorb

information from existing industrial IoT sensors and platforms.

to automatically diagnose failures. Uniquely, it also forecasts

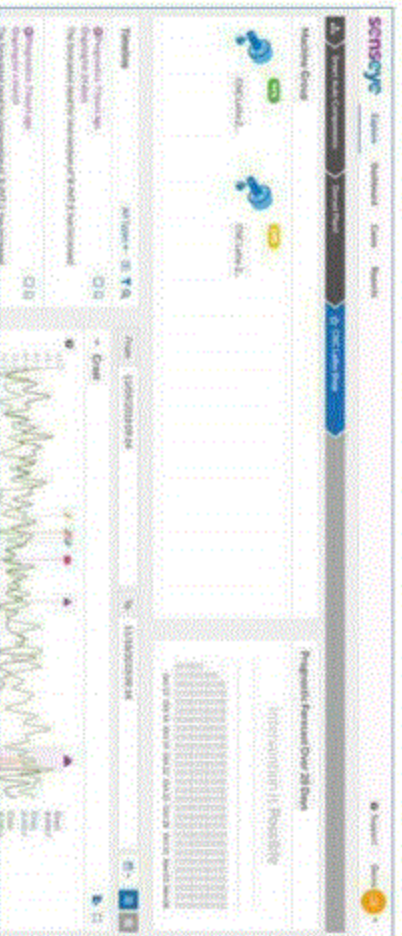
the remaining useful life of machinery.

Our solution can:

- monitor tens of thousands of assets, in real-time, without increasing effort for the end-user.
- maintain, as well as build, models of normal and abnormal machine behaviour, saving extraordinary effort and cost.

- reduce the cost of continuous condition monitoring by orders of magnitude
- reduce unplanned downtime by 50%, reduce the cost of continuous condition monitoring by orders of magnitude, and deliver a return on investment in under three months.

AI is frequently perceived as a threat, scaremongering in the press has caused unnecessary apprehension. These fears are misplaced. The AI we use helps talented engineers and maintenance professionals achieve more, and with greater focus. AI will not replace human experts – it will augment their abilities and allow them greater visibility of matters that require their attention and expertise. With AI taking care of the mundane, a symbiotic relationship paves the way for a new era of productivity.



Featured company

Signal

SIGNAL

Customers Businesses

Function Information & Regulation

Our platform uses AI to aggregate and analyse information in real time. Our technology translates and categorises content from 2.8 million digital, print and broadcast sources to deliver quick and easy access to relevant global data. We provide clients with the capability to monitor whatever subject matter they choose – organisations, people, events or topics – for myriad use cases including reputation management, regulatory compliance, business development and account management.

Our clients receive:

- real-time access to premium and exclusive content sources and datasets, in a single platform that saves time and money compared with subscribing to individual systems from multiple providers.
- our Insights dashboard, which enables clients to discover

emerging trends and track established topics that would be difficult, or impossible, to otherwise follow.

- a report builder, exportable graphs and email alerts, which simplify the process of sharing relevant news and insights clients' wider audiences, including their C-suites and external stakeholders.

Data is everywhere. The digital revolution has resulted in the 'datasphere' growing at an unprecedented rate. IDC predict there will be more than 163 zettabytes (163 trillion gigabytes) of data by 2025 – a tenfold increase from 2016. It's no longer humanly possible to process and extract salient information from this volume of data, and the challenge will continue to increase. Alongside this data deluge there is greater risk, competition, regulation and opportunity. As a result, AI will have a profound role to play in the corporate world.



Featured company

StoryStream



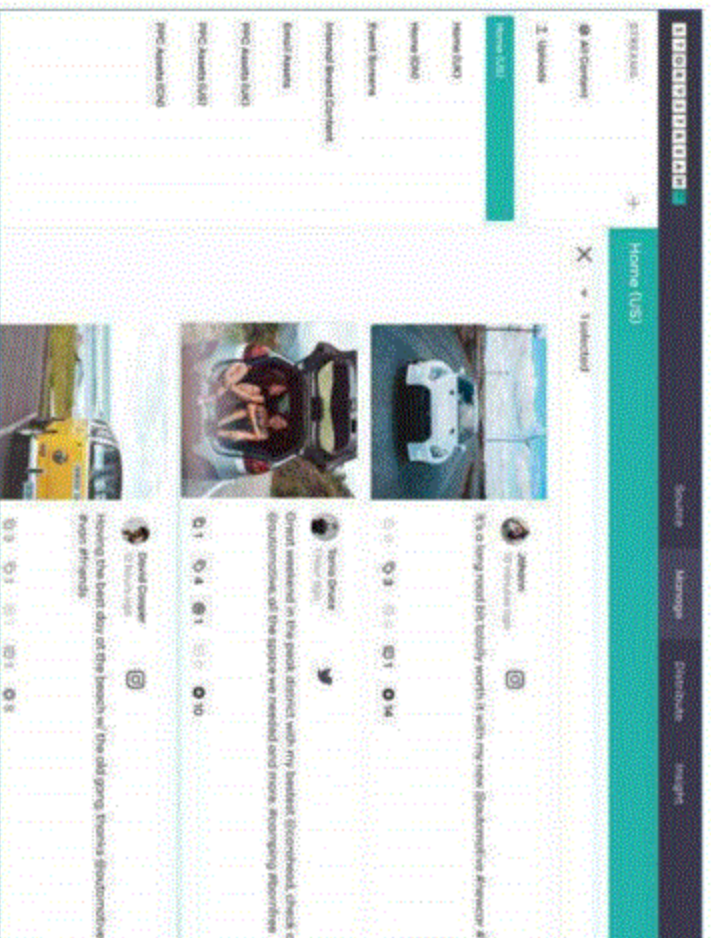
Customers Businesses Function Marketing

We offer a smart content platform for automotive brands.

Using patent-pending AI, our solution transforms the car-buying experience by automatically delivering more relevant, relevant content, at scale, to every customer touchpoint. Our platform enables better lead quality and increased conversion, and improved ROI by driving efficiency into content management across global teams.

- Using our solution, brands experience dramatic increases in marketing efficiency across their organisations and up to a 22% increase in website conversions; a 4.5x increase in customer engagement through targeted content; and a 5x ROI on their platform investment. Our proprietary AI, Aura, solves problems for automotive brands by:
 - enriching their content at scale by surfacing assets hidden, previously unavailable context.

- using image tagging and sorting to automatically create an organised content ecosystem, which underpins high quality marketing campaigns
 - automatically preparing and distributing content to brand stakeholders and customers, including intelligent content segmentation for personalised website experiences.
- AI enables intelligent context searching, which will enable marketing to become hyper-personalised. Marketers will readily be able to match individuals with the most appropriate content based on their buying behaviours. With the increasing power and performance of generative AI networks, we will begin to see artificial images used pervasively, compressing the time from concept to campaign and eliminating the need for expensive photo shoots or post-processing.



Featured company

Synthesisia

Customers Businesses

Core Technology Computer Vision


We offer responsible video synthesis technology to empower visual content creation. We remove the language barrier from video, by enabling 'native' translation that re-animates an actor's face to make it appear they speak a foreign language. In addition to synthesis technology, we are also developing tools to prevent malicious use of generative AI.

Our solution:

- delivers native translation of video content, which enables advertisers, celebrities and film & television producers to reach new, global audiences and establish new markets.
- enables content creators to develop different permutations of a single video, to create personalised adverts.
- allows celebrities to create new video content without having to be present on set.

AI will impact the media landscape in three ways:

- 'Native' translation of video will foster cultural exchange and deeper understanding – just as written translations have done for centuries.
- Hollywood will face increasing global competition, as high-end visual effects are democratised and the original language of video content becomes unimportant.
- Verification of digital identity and media will become as ubiquitous as secure socket layer (SSL) technology.



Chapter 8

The implications of AI

AI will have profound implications for companies and societies. AI will reshape sector value chains, enable new business models and accelerate cycles of creative destruction. While offering societies numerous benefits, AI poses risks of job displacement, increased inequality and the erosion of trust.

Summary

- AI's benefits can be abstracted to: innovation (new products and services); efficacy (perform tasks more effectively); velocity (complete tasks more quickly); and scalability (free activity from the constraints of human capacity). These benefits will have profound implications for consumers, companies and societies.
- By automating capabilities previously delivered by human professionals, AI will reduce the cost and increase the scalability of services, broadening global participation in markets including healthcare and transport.
- In multiple sectors including insurance, legal services and transport, AI will change where, and the extent to which, profits are available within a value chain.
- New commercial success factors – including ownership of large, private data sets and the ability to attract data scientists – will determine a company's success in the age of AI.
- New platforms, leaders, laggards and disruptors will emerge as the paradigm shift to AI causes shifts in companies' competitive positioning.
- AI, 'x-as-a-service' consumption, and subscription payment models will obsolesce select business models and offer new possibilities in sectors including transport and insurance.
- As AI gains adoption, the skills that companies seek, and companies' organisational structures, will change.
- By reducing the time required for process-driven work, AI will accelerate innovation. This will compress cycles of creative destruction, reducing the period of time for which all but select super-competitors maintain value.
- AI will provide profound benefits to societies, including: improved health; greater manufacturing and agricultural capability; broader access to professional services; more satisfying retail experiences; and greater convenience. AI also presents significant challenges and risks.
- AI-powered automation may displace jobs. AI will enable the automation of certain occupations that involve routine, in other occupations, AI will augment workers' activities. The short period of time in which select workers may be displaced could prevent those who lose their jobs from being rapidly reabsorbed into the workforce. Social dislocation, with political consequences, may result.
- Biased systems could increase inequality. Data used to train AI systems reflects historic biases, including those of gender and race. Biased AI systems could cause individuals economic loss, loss of opportunity and social stigmatisation.
- Artificial media may undermine trust. New AI techniques enable the creation of lifelike artificial media. While offering benefits, they enable convincing counterfeit videos. Artificial media will make it easy to harass and mislead individuals, and weaken societies by undermining trust.
- AI offers trade-offs between privacy and security. As AI-powered facial recognition advances, to what extent will citizens be willing to sacrifice privacy to detect crime?
- AI enables the high-tech surveillance state, with greater power for control. China is combining real-time recognition with social scoring to discourage undesirable activity.
- Autonomous weapons may increase conflict. The risk of 'killer robots' turning against their masters may be overstated. Less considered is the risk that conflict between nations may increase if the human costs of war are lower.

Recommendations

Executives

- Evaluate how the benefits unleashed by AI – innovation, efficacy, velocity and scalability – will impact your industry.
- Consider if AI can be used to reach new market participants and expand your addressable market.
- Assess the shifts in your industry value chain that will occur as adoption of AI grows.
- Evaluate the business model a disruptor might adopt in the age of AI. If freed from the “innovator’s dilemma”, what would the Netflix to your Blockbuster look like?
- Assess the extent to which your company is developing the commercial success factors required for the age of AI.
- Companies’ competitive positioning will change as adoption of AI increases. Develop an AI strategy to become a leader rather than a laggard.
- Evaluate the sustainability of your company’s skills and organisational design in light of changes AI will necessitate.
- Recognise the need for responsible stewardship. AI presents risks to society – including issues of job displacement, bias, and privacy. Develop rigorous ethical frameworks to govern the AI systems you develop and use.

Entrepreneurs

- Identify opportunities to take advantage of probable shifts in sector value chains that AI will cause.
- Develop initiatives that will take advantage of the new market participants and business models that AI will present.
- Identify weaknesses in incumbents’ competitive positioning that are likely to persist, or worsen, given their structure or strategy.
- Be mindful of the risks AI poses to society. Develop robust frameworks for ethical development and regulatory compliance. Explore Chapter 8 of our AI Playbook (www.mmcventures.com/research) for an actionable guide.

Investors

- Assess how the innovation, efficacy and scalability enabled by AI will impact your portfolio companies.
- Identify investment opportunities in sectors that will be transformed as a result of AI altering value chains and enabling new market participants.
- Evaluate opportunities to invest in companies structured around business models that will come of age as AI disrupts existing markets.
- Assess entrepreneurs’ awareness of AI’s ethical risks, their mitigation strategies and compliance with regulatory best practices.

Policy-makers

- Engage with experts in the field of AI bias to highlight the risks posed by prejudiced systems, create frameworks for best practice and highlight non-compliance.
- Engage the public in debate regarding the trade-off desired between privacy and AI-enabled security.
- Anticipate the proliferation of artificial media, and work with technology and media companies to support the creation of systems of trust.

AI will deliver innovation, efficacy, velocity and scalability

AI's value, from finding patterns in data more effectively to automating previously manual tasks, can be abstracted to four key benefits (Fig. 94):

Fig 94. AI offers innovation, efficacy, velocity and scalability

Benefit	Explanation	Examples
Innovation	New products and services.	<ul style="list-style-type: none">• Autonomous vehicles• Voice-controlled devices
Efficacy	Perform tasks more effectively.	<ul style="list-style-type: none">• Fraud detection• Customer segmentation
Velocity	Complete tasks more rapidly.	<ul style="list-style-type: none">• Legal document processing• Manufacturing process optimisation
Scalability	Extend capabilities to additional market participants.	<ul style="list-style-type: none">• Automated medical diagnosis• Automated executive assistants

Source: MNC Ventures

AI will have significant implications for markets and societies

Innovation, efficacy, velocity and scalability will have significant implications for economic systems, employees, consumers and society.

Below, we explain how AI will disrupt companies and markets by enabling:

1. **New market participants**
2. **Shifts in sector value chains**
3. **New business models**
4. **New commercial success factors**
5. **Changes in companies' competitive positioning**
6. **Shifts in skills and organisational design**
7. **Accelerated cycles of innovation.**

For societies, in addition to numerous benefits AI presents challenges and risks. Below, we describe how:

1. AI-powered automation may displace **jobs**
2. Biased systems could increase **inequality**
3. Artificial media will undermine **trust**
4. AI offers states greater **control** and presents trade-offs between privacy and security
5. Autonomous weapons may increase **conflict** between nations.

1. New market participants

By automating capabilities previously delivered by human professionals, AI will reduce the cost and increase the scalability of services, significantly broadening participation in select markets.

Today, access to sectors including healthcare and financial services is limited to subsets of the global population.

Medical diagnosis, for example, is inaccessible to people in developing economies and expensive for those in developed nations. Diagnosis has been undertaken by experienced professionals, whose training is time consuming and whose scalability is limited, inhibiting supply and increasing cost.

AI will provide automated diagnosis for a growing proportion of conditions. The marginal cost of diagnosing a patient with an AI algorithm will be nil. With smartphone adoption in developing economies increasing rapidly, from 37% in 2017 to an estimated 57% by 2020 (GSMA), barriers to access are also falling rapidly. By transferring the burden of diagnosis from people to software, global access to primary care will increase. Millions of additional individuals will benefit from primary care, while the market for providers of relevant and associated technologies will expand.

By automating capabilities previously delivered by human professionals, AI will reduce the cost and increase the scalability of services, broadening participation in select markets.

2. Shifts in sector value chains

In multiple sectors AI will change where, and the extent to which, profits are made within a value chain.

In the insurance sector, revenue from car insurance accounts for 42% of global insurance premiums (Autonomous Research).

42%
of revenue from global
insurance premiums come
from car insurance.

Source: Autonomous Research

As AI-powered autonomous vehicles gain adoption, the frequency of accidents will reduce – and with them, insurers' revenue.

UK car insurance premiums are expected to fall by as much as 63%, causing profits for insurers to fall by 81% (Autonomous Research). Insurers must anticipate and plan for a profound shift in their sector's value chain.

In the legal services sector, clients are increasingly aware, and less willing to pay, for deliverables that have not required the time or expertise of an experienced lawyer. In March 2017, Deutsche Bank announced that it will no longer pay City law firms for legal work undertaken by trainees and newly qualified lawyers. The automation enabled by AI will broaden the range of tasks that can be provided to clients at low cost. As clients expect greater use of AI, cost pressures on routine work will increase and value will shift further to high-end work.

In the transport sector, automotive finance provides 19%, on average, of car manufacturers' pre-tax profits (MMC Ventures). Large automotive finance companies, including Ford Motor Credit, Toyota Financial Services, Nissan Motor Acceptance Corp and Hyundai Motor Finance loan consumers money to buy new cars. As we describe next ("New business models"), private vehicle ownership will reduce as subscription-based services provide consumers with on-demand access to fleets of autonomous vehicles. Demand for, and value in, automotive finance for consumers is likely to decline.

3. New business models

AI, growth of 'as-a-service' consumption, and subscription payment models will obviate select business models and offer new possibilities in sectors including transport, insurance and healthcare.

The greatest impact of new corporate and consumer technologies is the new business models they enable, not the technical capabilities they provide.

In the transport sector, AI will transform the economic fabric of ownership and insurance. Cars are parked for an average of 96% of their lives (UTP Millennium Cities Database). Despite the cost and inefficiency of private car ownership, the model has been necessary to enable spontaneity, point-to-point convenience, comfort, privacy and security during travel.

An autonomous vehicle, summoned whenever required from a distributed fleet and used for the duration of a journey, will offer the same benefits while optimally utilising a fleet.

With the cost of the driver removed, and the cost of the vehicle and insurance divided over a greater volume of trips in a given period, the marginal cost of a journey will be lower. With growing use of transport-as-a-service subscription models, in which consumers pay a low monthly fee for on-demand access to a fleet of autonomous vehicles, private car ownership is likely to decline.

The impact on 'downstream' market participants will be as significant. The business models of local car dealerships, vehicle repair centres, petrol stations and charging centres will change as local ownership of private vehicles is displaced by large, managed fleets.

In the insurance sector, associated business models will be disrupted. The object of car insurance is likely to change, from a driver (who will play no role in an autonomous vehicle's operations) to the vehicle manufacturer or service provider. The immediate buyer of car insurance will also change.

From the end user to the manufacturer or service provider. (Ultimately, the fee will be repaid by the end user as a small component of their monthly subscription fee). Accordingly, insurers' business models in the automotive sector may shift from private policies to fleet-based agreements. Today, 67% of car insurance policies are personal, not commercial. This may fall to 40% (Autonomous Research).

In the transport sector, AI will transform the economic fabric of car ownership and insurance.

4. New commercial success factors

New commercial success factors will determine a company's ability to be successful in the age of AI.

A paradigm shift in technology offers companies new benefits while demanding new competences. Cloud computing, for example, offered flexibility, scalability, reduced capital expenditure and faster upgrade cycles. However, it demanded new diligence processes, different supplier relations and dynamics, internal competencies in change management and greater attention on security.

Success factors in the age of AI include:

- The vision to embrace AI and the organisational changes it requires;
- Ownership of large, non-public data sets to train and deploy market-leading AI algorithms;
- A willingness to evaluate the opportunities and risks of sharing training data with partners and competitors;
- The ability to attract, develop, retain and integrate data scientists within an organisation;
- The ability to form effective partnerships with best-of-breed third-party AI software and service providers;
- The ability to disengage AI partners effectively;
- A willingness to understand and respond to regulatory challenges posed by AI;
- A shift in mindset to the use of software that provides probabilistic instead of binary recommendations;
- The ability to manage workflow changes that result from the implementation of AI systems;
- The ability to manage challenges of organisational design and culture as AI augments, and in some cases replaces personnel.

5. Changes in companies' competitive positioning

New leaders, followers, laggards and disruptors will emerge as the paradigm shift to AI causes significant shifts in companies' competitive positioning.

Paradigm shifts in technology destabilise incumbents and enable new leaders to emerge. As adoption of cloud computing continues, for example, IT spend is being reallocated to cloud-native platforms (such as Amazon) and applications at the expense of incumbents.

AI will cause greater shifts as it alters value chains, enables new business models and demands different success factors from competitors. We expect 'Platforms', 'Disruptors', 'Leaders' and 'Laggards' to emerge.

New leaders, laggards, platforms and disruptors will emerge.

Among providers of AI:

Platforms – primarily Google, Amazon, IBM and Microsoft (GAIM) – provide the AI infrastructure, development environments and 'plug and play' AI services used by many developers and consumers of AI. With vast data sets, world-class AI teams and extensive resources, select GAIM vendors are well positioned to accrue value as platforms that support the provision of AI.

GAIM do not, however, have the data advantage, expertise or strategic desire to address the myriad domain-specific use cases required by businesses in sectors ranging from manufacturing, agriculture and education to retail, professional services and finance. This presents opportunities for disruptors.

Disruptors are early stage, AI-led software companies tackling business problems in a novel way using AI. For incumbents, disruptors are a double-edged sword. Disruptors will enable the enterprises, small- and medium-sized businesses that embrace them, while eroding the value of those that lack the foresight to do so. Select disruptors will become tomorrow's incumbents or be acquired by today's.

Among buyers of AI (today's enterprises, and small and medium-sized businesses):

Leaders will emerge in key industries, by: anticipating the shifts in value chains and business models caused by AI; taking advantage of their large, proprietary data sets to train and deploy AI algorithms; having the organisational ability to deploy AI effectively; and by having sufficient resources and reputation to attract high quality AI talent. Leaders will extend their competitive advantage and enjoy particular benefits:

1. In the 'data economy', economic returns will accrue disproportionately to companies that can extract value from information most effectively.
2. Data network effects create wider competitive moats. Larger volumes of training data enable better algorithms, which deliver better products and services, which win more customers, who provide more data. Leaders will benefit from data network effects that competitors will struggle to overcome.

Laggards are buyers that lack the will or organisational ability to use AI effectively. While some enterprises will lack the foresight to adapt, more will falter due to limited organisational capability. Laggards will move slowly to partner with disruptors or invest in their own AI teams, fail to take advantage of the extensive data sets and resources at their disposal, and struggle to attract AI talent. In the 'data economy', laggards will lose competitive advantage and market share significantly and rapidly.

New leaders will anticipate the shifts in value chains and new business models enabled by AI.

6. Shifts in skills and organisational design

As AI gains adoption the skills that companies seek, and companies' organisational structure, will change.

As companies vie for leadership in the AI era, companies will seek different personnel and change the organisational principles around which they are structured.

41% of companies are considering the impact of AI on future skill requirements (PwC). A mix shift to employing data scientists is likely. Data scientists extract meaning from data by collating, cleaning and processing data and then applying statistical techniques and AI algorithms. Companies' engagement with data scientists is limited today. For example, while the world's largest professional services and consulting firms average 5,000 to 15,000 in-house analytics professionals, we estimate that fewer than 8% of these are data scientists (MMC Venture). Some large companies have as few as 100 data scientists. Tomorrow's leaders are aggressively expanding their data science teams, recognising that time to market is key because of the potential for competitive advantage through data network effects (more data yields better algorithms, which provide improved products that attract more clients and data).




Fig 95. Cycles of innovation, adoption and consumption are compressing

While adjusting their mix of personnel, companies will alter their organisational design. Hiring for adaptability will be increasingly important, as the range of tasks supported or undertaken by AI systems increases. One in three companies are redesigning their organisational structures from traditional hierarchies to multi-disciplinary teams (Deloitte) to enable greater adaptability.

7. Accelerating cycles of innovation

By reducing the time required for process-driven work, AI will accelerate the pace of business and innovation. This may compress cycles of creative destruction, reducing the period of time for which all but a select number of super-competitors maintain value.

With several occupations, and numerous constituent activities, automated or augmented with AI, the speed at which tasks can be completed will increase. By accelerating the pace of business, AI is likely to shorten cycles of innovation, adoption and consumption that have been compressing since the 1950s (Fig 95).

Historically, accelerating cycles of innovation have reduced the period of time for which large companies retain value. In 1965, companies in the S&P500 stayed in the index for an average of 33 years (Innosight). By 1990, average longevity had narrowed to 20 years. By 2012, 18 years was typical. By 2026, average tenure in the S&P 500 is forecast to shrink to 14 years (Innosight). While reduced longevity in stock market indices arises partly due to technical factors, such as increasing merger and acquisition activity, creative destruction of incumbents has been accelerating. Faster cycles of disruption due to AI could reduce, further, large companies' ability to maintain value.

However, the dynamics of AI, and today's market leaders, may result in a divergence in longevity and the emergence of a small number of super-competitors. Three factors could lead to the emergence of super-competitors that maintain value for longer than companies in recent history.

First, AI offers network effects through data. Because training AI algorithms typically requires large volumes of data, companies with large, proprietary data sets can deliver more effective AI systems. Superior systems provide better results, which attract more customers, who bring additional data – creating a virtuous circle and powerful defensibility. Several of today's largest technology companies including Google, Amazon, Apple and Microsoft have vast consumer data sets inaccessible to disruptors.

Artificial media will make it easy to mislead – to harm individuals by ascribing to them words they have not said and actions they have not performed.

Second, today's leading technology companies are investing, and expanding, into emerging technologies and product categories more forcefully than many companies in the past. Leading technology companies are disrupting themselves. Google, a company conceived to index pages on the world wide web, has become a leader in autonomous vehicles and quantum computing. Amazon, a company that sold books online, is becoming a force in so many sectors that the Company is mentioned on 10% of all US company quarterly earnings calls (Reuters).

Third, select 21st century technology companies are consolidating power by expanding up, and down, the technology stack. Providers of cloud storage, such as Amazon and Microsoft, are layering ever-higher levels of functionality – such as AI and security – into the environments they provide. Technology leaders are also expanding down the technology stack. Google and Apple now develop their own microprocessors for competitive advantage in mobile and AI computing. By expanding up and down the technology stack, companies can consolidate control and customer spend.

The combination of data network effects, greater investment in emerging technologies and product categories, and expansion up and down the technology stack may enable a small number of super-competitors to capture and maintain economic influence for a longer period of time than has been possible in recent history – amidst a broader bifurcation in corporate longevity.

AI offers benefits and risks to societies

AI will deliver numerous, profound benefits for societies. They include: accelerated cycles of innovation; broader access to better, less expensive healthcare; increased manufacturing capability and agricultural productivity; enhanced mobility with fewer accidents; improved management of financial assets and risk; broader access to lower-cost professional services; more efficient and satisfying retail experiences; and greater day-to-day convenience.

AI also presents significant challenges and risks.

Below, we describe how:

1. AI-powered automation may displace jobs;
2. biased systems could increase inequality;
3. artificial media will undermine trust;
4. AI offers states greater control and presents trade-offs between privacy and security; and
5. autonomous weapons may increase conflict between nations.

Increasingly, AI is enabling divergent futures. The extent to which risks crystallise will depend upon the choices and actions of citizens, organisations, companies and governments.

1. AI-powered automation may displace jobs

Job displacement is a significant risk associated with the proliferation of AI. AI will directly enable the automation of several occupations that involve routine and repetition – from truck-driving to telemarketing. Truck driving comprises 3.6 million jobs in the US (American Trucking Association). In many other occupations, AI will augment and then displace some workers in more complex roles, while reducing the need for additional workers to be hired as companies expand. In approximately 60% of occupations, at least 30% of consultant activities are technically automatable by adapting currently proven AI technologies (McKinsey Global Institute).

Analysis of UK census data since 1871 shows that historically, contracting employment in agriculture and manufacturing – a result, in part, of automation – have been more than offset by rapid growth in the caring, creative, technology and business service sectors (Deloitte).

Greater automation of both manual and business service roles, however, may concentrate employment further in occupations resistant to automation, including care work and teaching. Whether or not, over time, AI creates more jobs than it destroys, the short time frame in which a large number of workers could be displaced, coupled with a reduction in the availability of similar roles, could prevent those who lose their jobs from being rapidly re-absorbed into the workforce. Social dislocation, with political consequences, may result.

2. Biased systems could increase inequality

Theoretically, AI has the potential to free decision-making from human bias by finding objective patterns in large data sets. However, AI systems typically learn by processing training data. Available data sets frequently reflect systemic historic biases, including those of gender and race.

The results from 'word embedding', an AI technique used to interpret written and spoken language, are an example. Word embedding creates mathematical representations of language. The meaning of words are abstracted to a set of numbers based on the words that frequently appear near to them. However, when trained on the Common Crawl data set (a 145-terabyte collection of data taken from material published online, the word 'woman' is closely associated with occupations in the humanities and the home, while 'man' is associated closely with science and technology professions (Caliskan, Bryson and Narayanan).

Lack of diversity among AI development teams is compounding the problem. Groups representing majorities in the population are less likely to notice that data regarding minorities is lacking in training data they use. In a popular data set for training facial recognition systems, over 75% of faces are male and 80% are lighter-skinned (Buolamwini, Gebru).

Inadequate or imbalanced training data are causing AI systems to perform poorly and problematically, particularly when serving minorities. For example, AI-powered facial recognition systems that offer gender classification misgender just 1% of lighter-skinned males – but up to 7% of lighter-skinned females, 12% of darker-skinned males and 35% of darker-skinned females (Fig. 96) (Buolamwini, Gebru).

Fig 96. AI-powered facial recognition systems misgender 1% of lighter-skinned males but 35% of darker-skinned females



Source: Buolamwini, A.I.T. (Healy lab), via The New York Times

Algorithms will make decisions that have significant ramifications for individuals' lives, in a growing range of domains from recruitment to credit. If bias is not recognised and removed from AI systems, individuals will suffer economic loss, loss of opportunity and social stigmatisation (Fig. 97).

"If we fail to make ethical and inclusive AI, we risk losing gains made in civil rights and gender equity under the guise of machine neutrality." (Joy Buolamwini).

"If we fail to make ethical and inclusive AI, we risk losing gains made in civil rights and gender equity under the guise of machine neutrality."

Joy Buolamwini

"There is a battle going on for fairness, inclusion and justice in the digital world." (Darren Walker, via The New York Times).


To avoid 'automating inequality', developers can:

- recognise the challenge, as a starting point for action;
- develop diverse teams that reflect the communities they serve;
- create balanced, representative data sets;
- deploy ethics and testing frameworks for system validation.

Fig 97. There are potential harms from algorithmic decision-making

INDIVIDUAL HARMS		COLLECTIVE SOCIAL HARMS
ILLEGAL DISCRIMINATION	UNFAIR PRACTICES	
HIRING	EMPLOYMENT	
INSURANCE & SOCIAL BENEFITS		
HOUSING		
EDUCATION		
CREDIT		
DIFFERENTIAL PRICES OF GOODS	ECONOMIC LOSS	
LOSS OF LIBERTY		
INCREASED SURVEILLANCE		
STEREOTYPE REINFORCEMENT		
DIGNITARY HARMS	SOCIAL STIGMATISATION	

Fig 99. Given video of former President Obama, researchers synthesised photorealistic, new lip-synched video



Source: Sawjanbhai, Seitz and Kerenhucher Shalerman

GANs will progress from synthesising individuals to scenes. Footage of individuals and events will be generated, or altered, with little cost and effort, to create ‘fake news 2.0’ for political purposes or counterfeit evidence in criminal cases. As smartphones are used to record high-definition video, and videoconferencing solutions such as Skype and Facetime are used pervasively, source material is becoming plentiful.

The proliferation of artificial media poses immediate and secondary risks. In the short term, artificial media will make it easy to mislead – to damage individuals by ascribing to them words they have not said and actions they have not performed.

In the longer term, the rise of artificial media will undermine trust. Positively, citizens will learn to question whether the media they see is authentic. However, if any media can be counterfeited, all media is open to challenge. What can be believed? Adversaries have recognised that sowing doubt and confusion to divide populations and inhibit collective action is frequently more powerful than direct action over the long

term. In *Nineteen Eighty-Four*, the dystopian novel by George Orwell in which a ruling party persecutes independent thinking, citizens are taught to ignore what they see and hear. “The party told you to reject the evidence of your eyes and ears. It was their final, most essential command” (*Nineteen Eighty-Four*, George Orwell). In the decade ahead, as the unreal becomes real, society will grapple with challenges of truth and trust.

In the decade ahead, society will grapple with challenges of truth and trust.

4. AI offers states greater control and presents trade-offs between privacy and security

In the age of AI, citizens and governments must re-evaluate the balance between security and privacy they desire – while states could enjoy greater powers of social control.

AI-powered facial recognition systems offer unprecedented capability. Technical maturation coincides with the proliferation of high-resolution cameras. Every smartphone owner carries a camera in their pocket. Over 1.85 million CCTV cameras were in place in the UK as early as 2011; on average, a citizen is captured on CCTV an estimated 68 times per day (Cheshire Constabulary Camera Survey). To what extent will citizens and governments be willing to sacrifice anonymity and privacy to prevent and detect crime?

In the age of AI, citizens and governments must re-evaluate the balance between security and privacy they desire.

Further, the combination of AI and real-time analytics is enabling the high-tech surveillance state, with greater capacity for social control. With increasing accuracy, AI-powered gait analysis can recognise individuals from their shape and movement – even if their faces are hidden. “You don’t need people’s cooperation for us to be able to recognise their identity” (Huang Yongzheng, Weirix, via the Associated Press). China intends to combine real-time recognition with social scoring, to rate citizens according to their behaviour and habits. Individuals with undesirable behaviour may be inhibited from travelling, suffer reduced internet connectivity, be penalised when applying for government roles and be impeded from placing their children in desired schools.

5. Autonomous weapons may increase conflict between nations

Weapon systems have incorporated a degree of autonomy for decades. The Phalanx CWS weapon system, for example, defends ships in 20 countries’ navies from missile attacks. The Phalanx combines a 20mm rotating Vulcan cannon with an automated system to interpret radar data, decide whether a target is a threat and engage it.

However, the combination of AI-powered computer vision systems, AI-based decision-making algorithms and improved robotics are enabling humanoid and aerial drones with greater capability and autonomy. The risk of ‘killer robots’ turning against their masters may be overstated. Less considered is the possibility that conflict between nations may increase if the human costs of war are lower. A country that thinks twice about sending young people into conflict may be more adventurous if the only assets in harm’s way are equipment,

States could enjoy greater powers of social control.



Notes

100

