Behind the technology

Five key trends in Artificial Intelligence

Contents

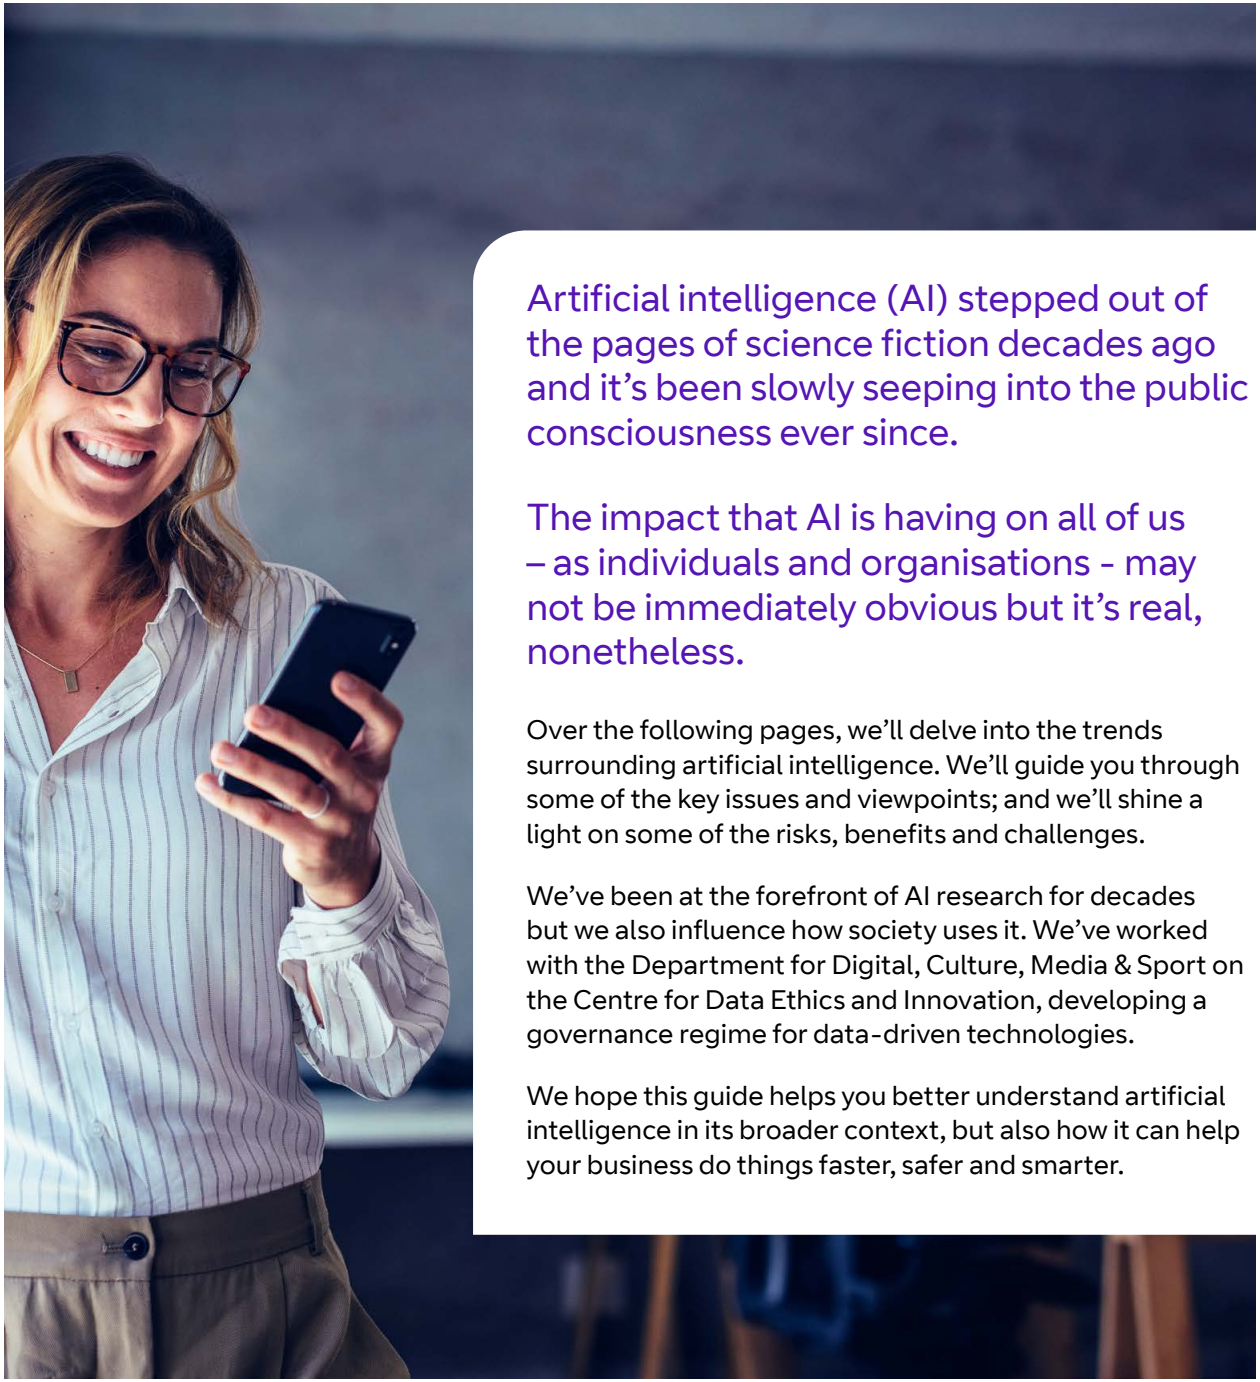

Artificial intelligence (AI) stepped out of the pages of science fiction decades ago and it's been slowly seeping into the public consciousness ever since.

The impact that AI is having on all of us – as individuals and organisations - may not be immediately obvious but it's real, nonetheless.

Over the following pages, we'll delve into the trends surrounding artificial intelligence. We'll guide you through some of the key issues and viewpoints; and we'll shine a light on some of the risks, benefits and challenges.

We've been at the forefront of AI research for decades but we also influence how society uses it. We've worked with the Department for Digital, Culture, Media & Sport on the Centre for Data Ethics and Innovation, developing a governance regime for data-driven technologies.

We hope this guide helps you better understand artificial intelligence in its broader context, but also how it can help your business do things faster, safer and smarter.

1. Filter bubbles

We don’t just use professional social networks like LinkedIn to build our contacts, we also use them to find new jobs. Using natural language processing and machine learning, AI job recommendation engines read and understand our profile, posts, and conversations to gauge the kind of roles we may be interested in.

They rank and prioritise the stories in our news feed by using machine learning algorithms to analyse what we’re likely to be interested in based on our past choices, and those of people who click on similar stories

These engines use a more advanced form of machine learning, known as deep learning, to process large amounts of data relating to people’s online behaviour. It’s not just the news we read that reveals our preferences – it’s also our web-browsing habits and what we post and click on social networks and messaging systems.

This activity is shared with providers in accordance with privacy policies we agree to when signing up to services.

By aggregating and analysing vast swathes of such data gathered from many people, deep learning systems are able to identify unusual patterns and connections in the data that make these recommendation engines ever more accurate.

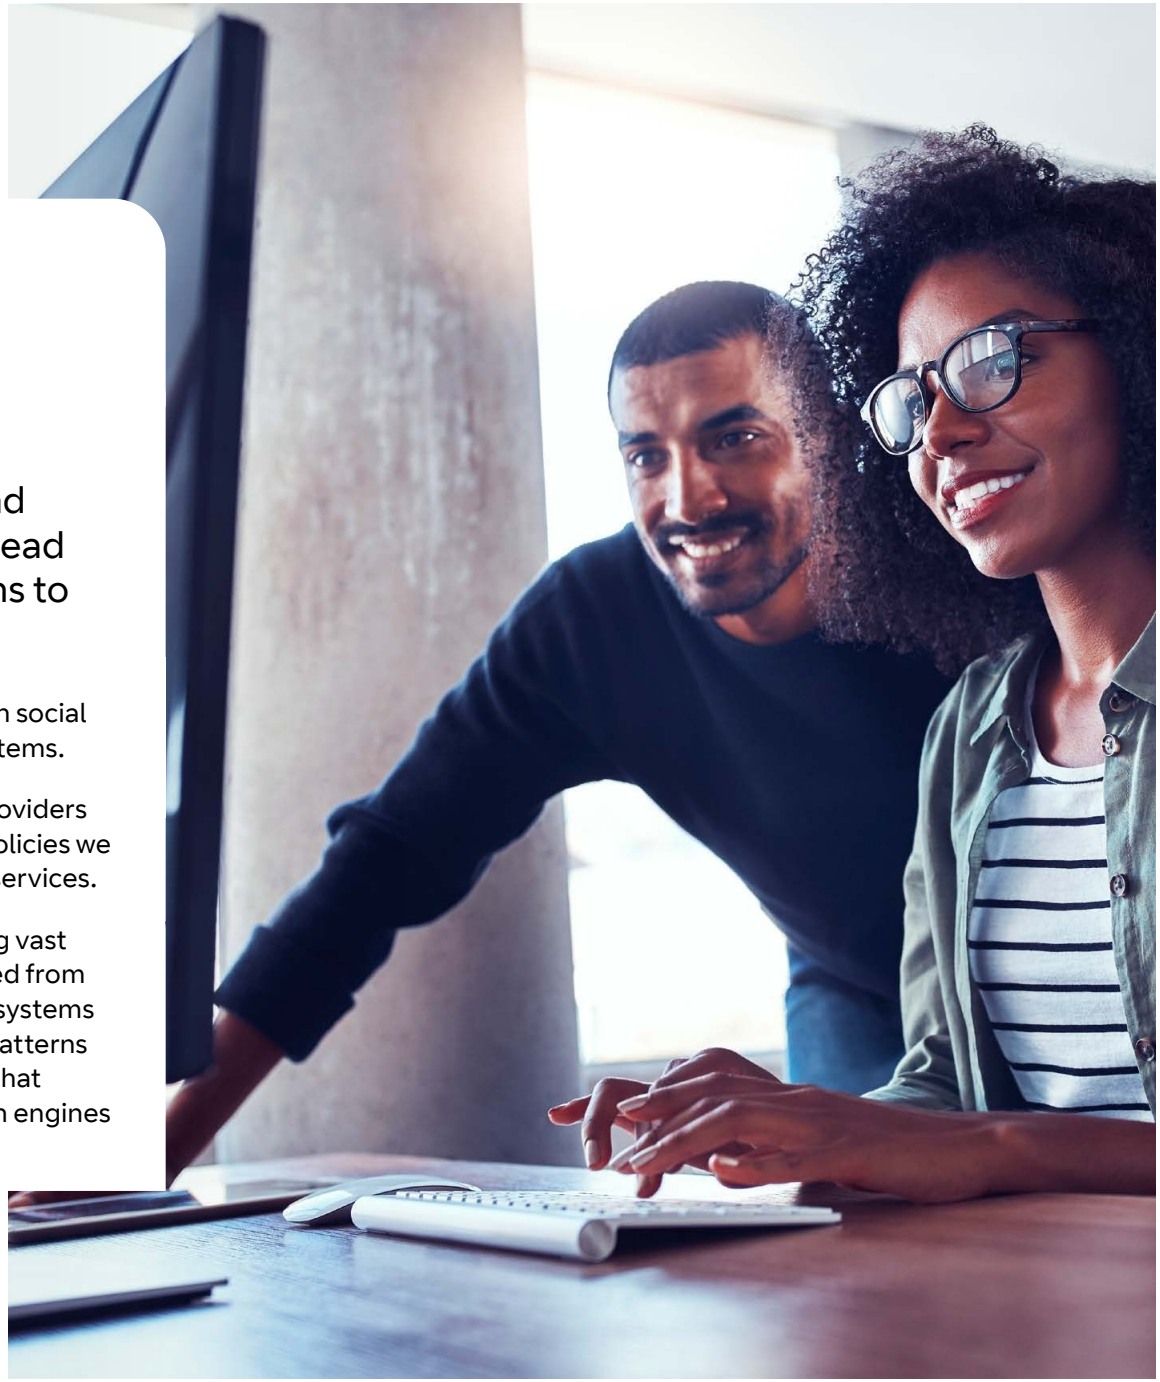

While this can be extremely helpful to people, companies should be mindful of processing personal data in this way.

Under the EU General Data Protection Regulation (the GDPR), organisations need to ensure that the processing of personal information is done under one of six legal grounds, of which consent is one.

In addition, many have raised concerns that such tools are creating, what prominent internet activist Eli Pariser calls 'filter bubbles' - the phenomenon by which people are increasingly fed stories that only reflect their own viewpoint, leading them to gain an ever more biased perspective on the world.

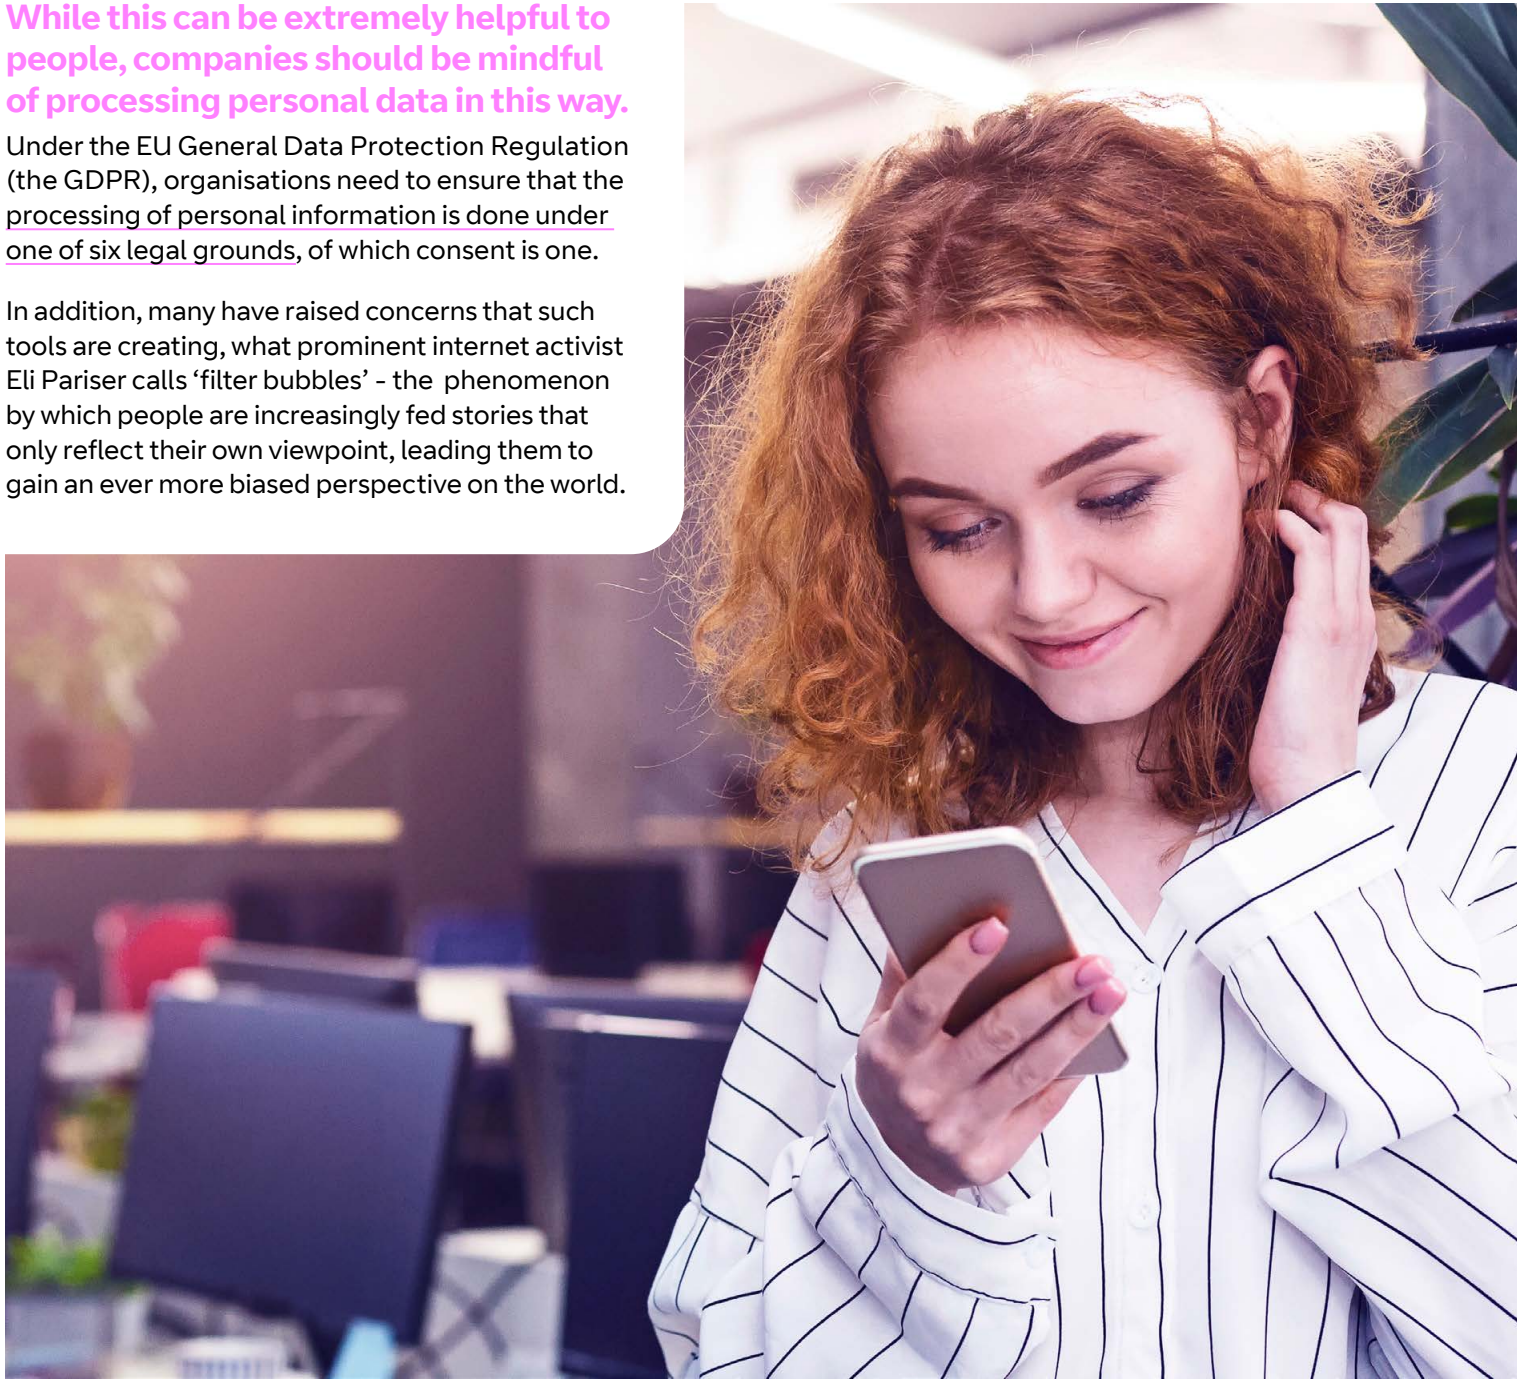

But this isn't all down to AI.

It's largely a consequence of people unwittingly creating their own echo chambers through who they choose to follow on social media (we tend to select people who reflect our own views), which is reinforced by the way that news and social platforms prioritise their content.

In fact, the big social media companies are beginning to look at ways in which AI can help counteract this effect by flagging up or rooting out fake or extremist stories and introducing more serendipity into people's news feeds.

2. Augmented Intelligence

Virtual agents, better known as chatbots, use natural language processing and machine learning to understand freeform text or voice data. They can act on any requests they detect and send appropriate, natural-language responses to people.

However, while chatbots can be effectively developed to carry out tasks with limited scope – such as arranging meetings or answering simple customer questions – it will be some years before they can understand and respond accurately to general conversation. In any case, it's important that companies make clear when users engage with chatbots rather than humans to avoid any confusion.

As AIs evolve and improve, certain jobs will very likely be displaced, particularly menial or repetitive tasks. Klaus Schwab, founder and executive chairman of the World Economic Forum, said "This will give rise to a job market increasingly segregated into 'low-skill/low-pay' and 'high-skill/high-pay' segments, which in turn will lead to an increase in social tensions."

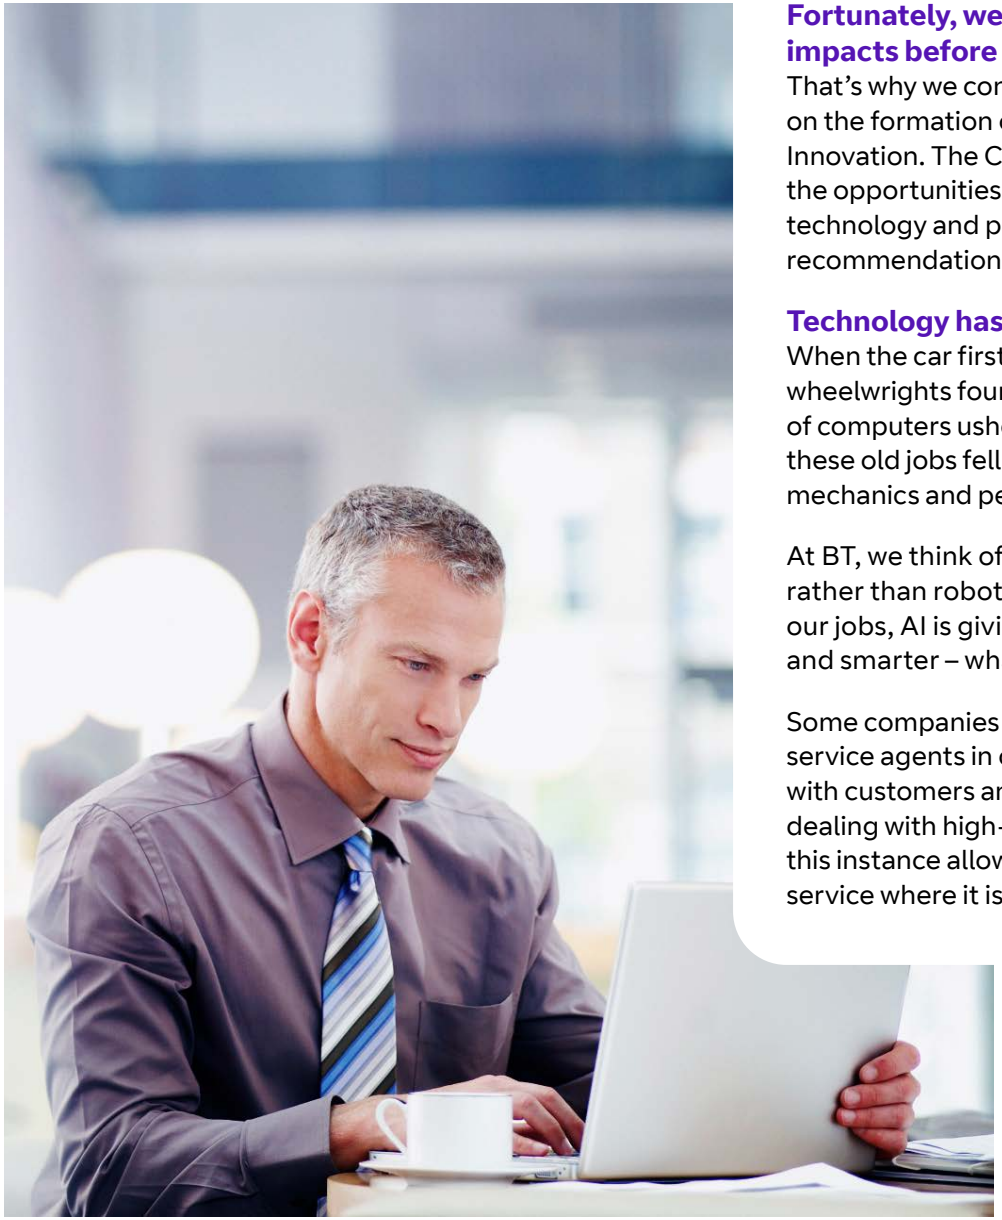

Fortunately, we have time to address these potential impacts before they hit.
That's why we contributed to the Government consultation on the formation of the Centre for Data Ethics and Innovation. The CDEI aims 'to analyse and anticipate the opportunities and risks posed by data-driven technology and put forward practical and evidence-based recommendations to address them'.

Technology has always driven societal change.
When the car first drove horse-drawn vehicles off the road, wheelwrights found themselves out of a job. The advent of computers ushered in the decline of secretaries. But as these old jobs fell away, new roles replaced them, like car mechanics and personal assistants.

At BT, we think of artificial intelligence in terms of 'cobots' rather than robots. In other words, rather than taking over our jobs, AI is giving us new powers to work faster, safer and smarter – what we call 'augmented intelligence'.

Some companies use chatbots to support their customer service agents in contact centres by observing interactions with customers and stepping in to help where they can. By dealing with high-volume, routine queries, AI chatbots in this instance allow human agents to focus on giving human service where it is most needed.

3. AI in the driving seat

Advances in machine learning technology, combined with ever more accurate sensors, cameras and position-tracking, are allowing cars to accurately ‘see’ other vehicles and obstacles, manoeuvre with high precision, and predict (and prevent) accidents.

Although automotive manufacturers are pushing hard to perfect **autonomous cars** – and have even introduced pilots in several cities across the world – the technology still has many hurdles to overcome before it becomes mainstream, human acceptance among them.

Bad weather or dirt can obscure cameras and sensors, positioning and location systems can fail, network connections can be unreliable, and the systems can be confused by unexpected events on the road or difficult terrain. There has already been one high-profile death caused by someone taking their hands off the wheel while using an autonomous car’s autopilot feature.

But as we move towards highly-networked smart cities – with vehicles, buildings and roadside beacons providing a constant stream of data that’s being analysed in real-time – AI technologies should help to improve traffic flows, pollution levels and road safety irrespective of whether or not we’re all in autonomous vehicles.

We’ll shortly launch our first connected vehicle offering, BT Auto Mate, which will provide real-time driver, vehicle health and tracking data of the kind that will be needed to fuel the AI-powered smart cities of the future. But for this vision to be realised, we need data to flow smoothly, speedily and securely at all times.

We’re responsible for the backbone of the UK internet.

It’s a responsibility we take seriously. That’s why we’re actively developing and using advanced AI technologies, making sure the network is ready and resilient enough to handle the explosion of real-time data traffic that these technologies will ignite.

Network security and resilience are our top priorities for AI. And as we train and deploy increasingly advanced **machine learning** algorithms to mitigate the risks of network failures and malicious attacks, our monitoring and defences are growing ever more capable and intelligent.

4. Explainable AI

Biometric face-recognition systems can be incredibly useful to businesses. Machine learning technology tracks faces on camera and compares them with those in its database You can do away with physical passes to access a building or a department if the ‘building’ itself recognises you and knows whether you’re authorised to enter.

But the use of such systems is controversial, throwing up difficult questions and ethical dilemmas.

In China, for example, the government’s plans to develop a facial recognition system for its 1.3 billion population in light of its proposed ‘social credit system’ raises concerns around wide-ranging surveillance and citizen privacy.

It’s important that companies and organisations only introduce modern technologies such as facial recognition after first consulting with people who’ll be affected by the technology: employees, service users, the local community and so on.

So in the case of deploying facial recognition in the workplace, for instance, you may want to give employees the option to continue using their old-style passes if they prefer.

HR departments can process applications quickly with AI CV-screening technology.

Advanced **machine learning** algorithms read and understand candidates’ applications and CVs, mining them for evidence of particular skills or traits. Some of the more advanced systems can also design personalised interview questions. And some companies are already developing the next step: chatbots that conduct preliminary interviews.

But just because you ‘can’ do it doesn’t necessarily mean that you ‘should’ do it.

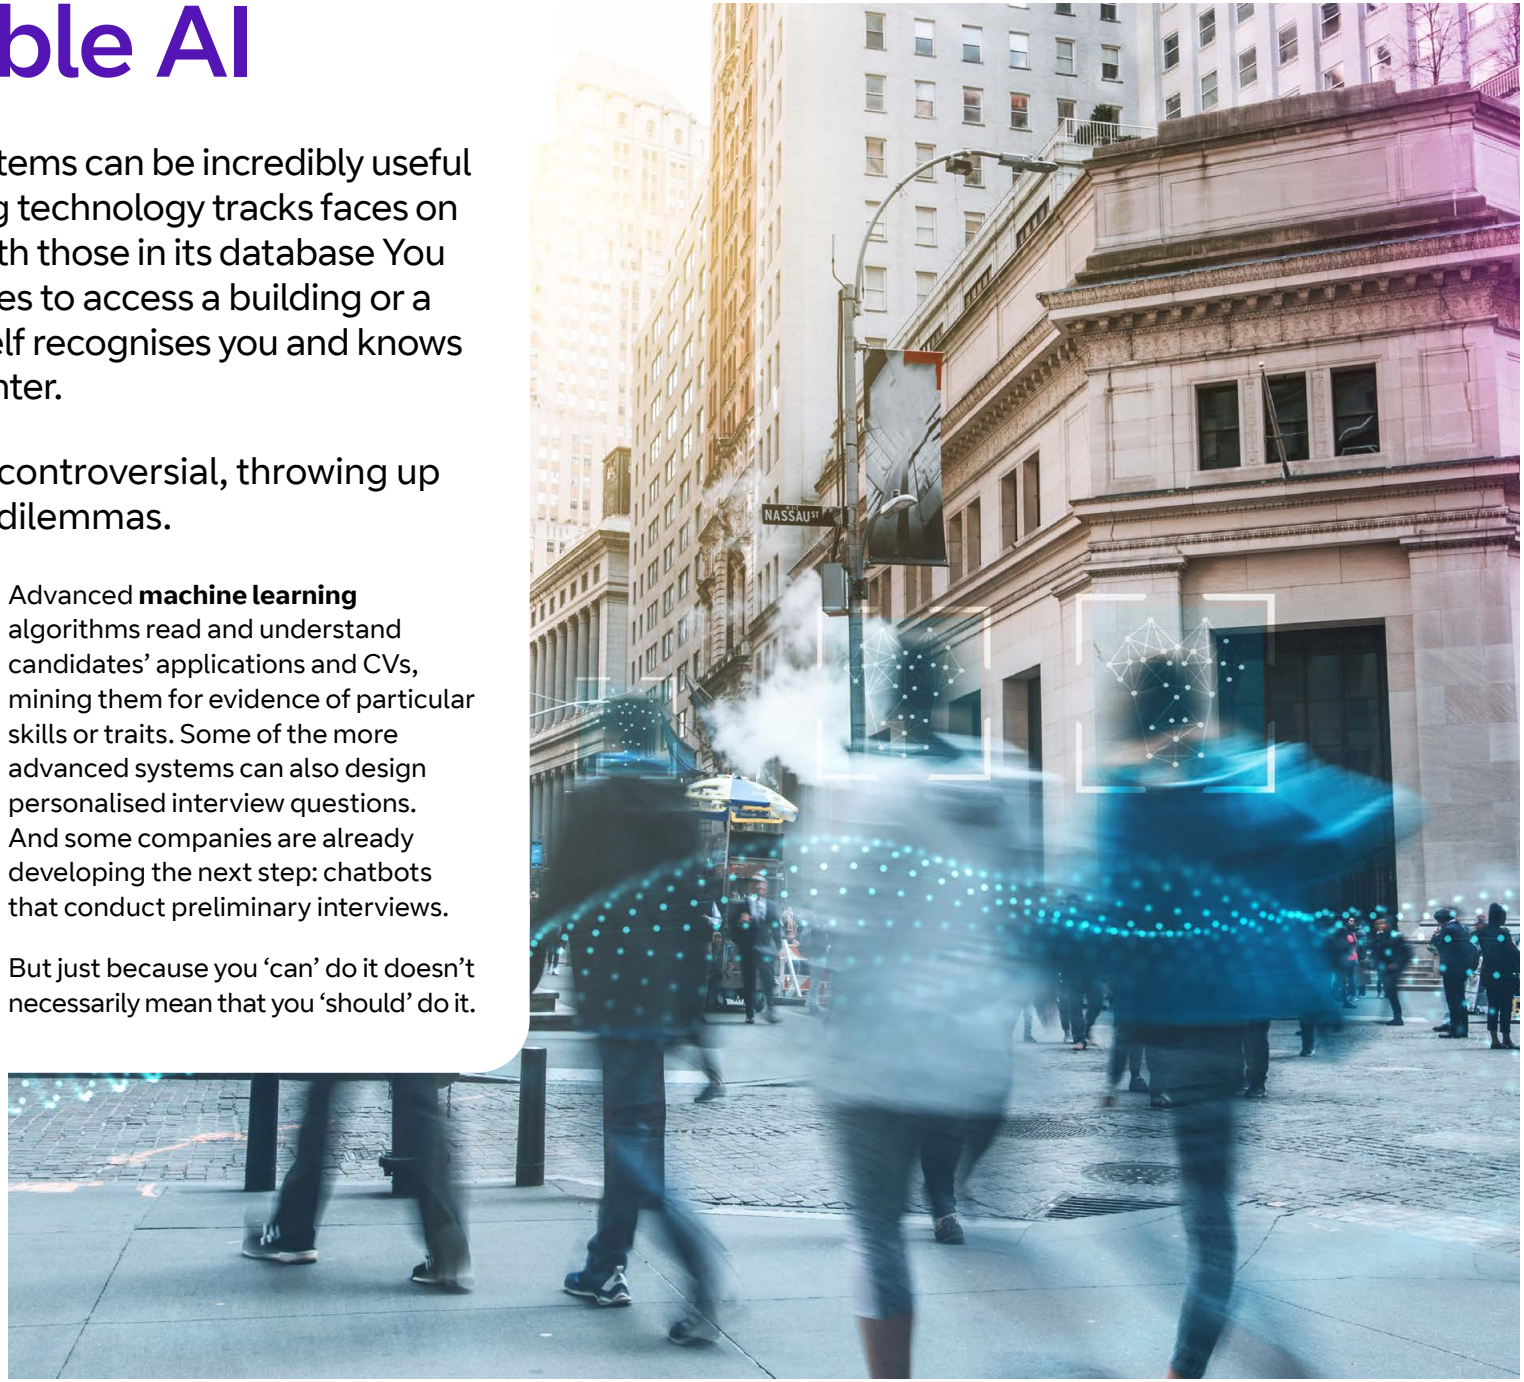

Using AI to make decisions over people's lives – such as whether they should be considered for a job – raises a number of ethical questions. For example, is it acceptable for firms to use AIs to trawl through candidates' online social media histories in order to assess their personality and dig up old misdemeanours?

Or to empower AI algorithms to reject candidates without anyone knowing precisely how they arrived at their judgments? One company recently withdrew a CV-processing AI system because it was biased against women even when gender information was left out of the CV. The problem with systems like that is that they are trained with historical data which always reflects historical bias.

As Dr Sandra Wachter, a lawyer and research fellow in data ethics at the Oxford Internet Institute at the University of Oxford, notes that AI algorithms can be 'highly complex and opaque'. She contends it is dangerous to deploy such systems when they are making decisions that affect people's lives, such as selecting and rejecting job applicants.

Conversely, others believe AI tools can actually make recruitment fairer, by eliminating human bias from the process, if the AI system is used to double-check a human decision.

Academics, governments, regulators and industry are busy tackling these questions and under certain legislation, including the EU's General Data Protection Regulation (GDPR), businesses must respect individual rights related to automated decision making and profiling, identifying when such processing takes place and allowing individuals to request human intervention or challenge a decision.

At BT we believe it's critical to build mechanisms into AI systems that reveal how they have arrived at their decisions. Our Applied Research division is busy developing what we call 'explainable AI' to do just that.

Detlef Nauck, our chief research scientist, explains: 'Although this research is still in its infancy, we see it as critical to the future acceptance and applicability of advanced AI technologies in business.'

5. AI and cybersecurity

One of the first commercial uses of AI was in the financial services industry. You know when your bank blocks a purchase because they think it's suspicious? That's fraud-detection AI in action.

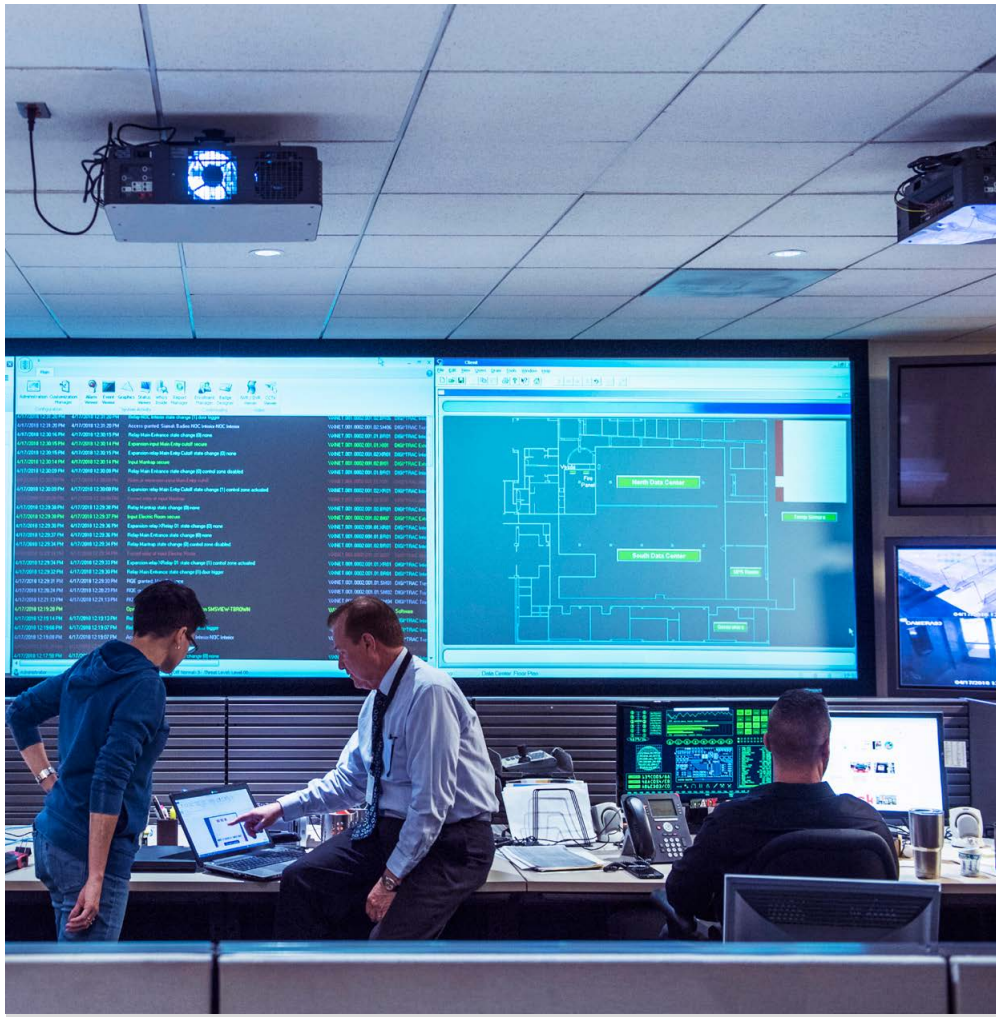

Machine learning and **neural network algorithms** examine the vast stream of transactions that banks process, looking for unusual purchasing patterns that could indicate fraudulent activity.

AI technology is now a standard element of security in all organisations

Having been trained using mountains of historical transaction data, these algorithms understand to a very fine degree what constitutes typical customer behaviour and what might signal criminal activity.

This enables them to flag up potentially suspect transactions almost immediately.

Today, AI technology has percolated through to general security solutions which can help all organisations. Security is our prime focus for AI. Technology such as our cyber security platform is helping to protect both our networks and those of our customers.

Machine learning and **deep learning** algorithms detect anomalies on the network and Saturn, our powerful in-house visual analytics system, alerts human experts to potentially problematic activity.

It also allows them to drill down and explore the data in a rich, visual way so they can rapidly gain a more nuanced understanding of what's happening and deal with any problems before they have the chance to cause any significant damage.

At the 2012 London Olympics, we used automated defences to prevent systems being brought down by DDOS attacks.

And as hacking and state-sponsored cyberterrorism techniques become ever more sophisticated (we're seeing a growing number of attackers deploying 'adversarial AI' against our own AI technologies) we're not standing still. We're constantly working on innovations in detection technology so we can counter any malicious attempt to interfere with our AI systems.

The threat landscape is so complex and fast-moving that it's critical to stay ahead of the game. We use AI to put us on the front foot, rather than having to be reactive.

Glossary

Adversarial AI

A technique used by hackers to trick machine learning systems through the use of slightly modified external data inputs. Although an image or other data may appear legitimate to a human observer, subtle modifications can cause it to be misclassified by the AI. This can corrupt the effectiveness of an organisation's security defences, allowing attackers to gain entry.

Artificial intelligence (AI)

A collection of techniques that enable computers to mimic human intelligence processes such as learning, reasoning and self-improvement. These can be further categorised as weak (or narrow) AI, which refers to systems built and trained to simulate a particular task, and strong AI – systems with general cognitive abilities that can find solutions to unfamiliar tasks. All of the AI technologies in use within business today are weak AI. Strong AI systems do not yet exist.

Artificial neural network (ANN)

See neural network.

Augmented intelligence

An alternative term for AI that emphasises its assistive qualities rather than its artificiality. A number of organisations are adopting the term in a bid to reframe negative public perceptions of AI that have often been informed by dystopian science fiction. The aim of augmented intelligence is not to replace human agency, but to support decision making and to automate menial steps allowing the human user to concentrate on the important aspects of the task.

Autonomous vehicles (AVs)

Self-driving cars or other vehicles that use techniques such as deep learning and machine vision to obey the rules of the road, avoid obstacles and deal with emergency situations such as pedestrians walking in front of them or erratic driving by other road users.

Biometrics

The measurement and analysis of people's unique characteristics, used primarily for access control and identification. These characteristics could be physical or behavioural – a fingerprint or retina scan, the proportions of a person's face, the unique timbre of their voice, their gait, etc. Combined with AI techniques such as **computer vision** and **machine learning**, biometric technology can today perform speedy identification, for example.

Chatbot/virtual assistant

A computer program that is able to engage in a dialogue with people using speech and/or text. While one enduring goal of AI is to create a convincing general chatbot that can converse on any topic as naturally as a human (see **Turing Test**), we're not there yet. Commercial chatbots in use today are typically limited to specialised conversations, such as dealing with sales and customer support enquiries, dealing with earlystage job applicants or acting as a phone operator.

Cognitive analytics/computing

While there is no widely agreed definition in academia or industry, the term is often used by AI companies for distinction. Typically, cognitive computing is associated with machine learning, reasoning, natural language processing, speech recognition, vision, human-computer interaction amongst other technologies.

One objective of cognitive computing is to create more accurate models of how the human brain/mind senses, reasons, and responds to stimulus.

Computer vision/machine vision

The ability of computers to interpret visual data such as a camera feed or video file. For example, a system might identify objects from a moving image, spot defective products on a production line or pick out individual faces from CCTV footage.

CV-screening software

See resumé-screening software.

Data analytics

The scientific process of turning data into insight for making better decisions. The methods used in data analytics typically come from statistics and machine learning.

DDOS attack

An advanced form of machine learning based on neural networks. Deep learning systems are especially suited to applications like image analysis and speech or natural language processing. They rely on large amounts of labelled data and require a lot of computational power to train them. They have also been successfully used in learning to play games like Go, chess or poker where they have been trained through reinforcement learning.

Evolutionary algorithm

A computer program that mimics biology to improve over successive generations by mutation, recombination and selection. This can either be autonomous or have their evolution guided by a human operator. The technique is commonly used to optimise systems, processes, products or problem-solving.

Evolutionary optimisation

See evolutionary algorithm.

Expert system

A computer system that replicates the decision-making ability of a human expert, solving tricky problems.

Expert systems originally became popular in the 1980s and were one of the earliest successful uses of AI in business. They are designed to solve complex problems by reasoning, using databases of knowledge and a series of pre-configured rules.

Explainable AI

An artificial intelligence whose actions and decisions can be transparently reviewed and understood by humans.

Fuzzy logic

Traditional logic represents data as either true or false. Fuzzy logic, invented in the 1960s, introduced the idea of vagueness that allowed for more nuanced interpretations of data, more akin to human reasoning.

It was used extensively in consumer electronics and you might find it in your camera (image stabilisation) or in your washing machine.

Heuristics

Techniques designed for solving complex problems more quickly by finding an approximate solution. A common use is in computer security systems that scan for unknown viruses by looking for patterns of activity that are similar to those exhibited by other viruses.

Intelligent agent

Computer security systems that use AI techniques like **machine learning** and **deep learning** to identify suspicious or anomalous activity on the network. The sheer number of threats and sprawling complexity of today's networks means IT departments can no longer rely on human or rules-based monitoring to secure an organisation's systems. From detecting potential fraud or suspicious activity on the network to preventing **DDOS attacks**,

AI techniques are increasingly key.

Machine learning

An AI technique by which computers can learn how to perform a specific task rather than being explicitly programmed to do it.

This is typically accomplished by presenting the system with lots of correct and incorrect examples of whatever it is you're training it to recognise. For example, an AI quality-control system might be fed many images of defective products and non-defective products, being told which are which, until it develops the ability to detect the difference to a very fine degree on its own. Machine learning is the basis of many of the AI systems in use within business today. See also **deep learning**.

Machine vision

See **computer vision.**

Natural language processing (NLP)

A collection of techniques, both AI-based and non-AI based, by which computers are able to extract meaning and context from text or speech in natural language, and/or communicate with people using natural language. These techniques are becoming increasingly sophisticated and are commonly used in **chatbots** and virtual assistants, for example.

Network intelligence

The ability to understand what traffic is passing across a network, and how it behaves, in real time.

This might involve various AI techniques to identify patterns of communication, to aid the routing and prioritisation of traffic for optimal performance, or to spot network security issues.

Neural networks

A neural network, or artificial neural network (ANN), is a framework for **machine learning** algorithms that is based on the idea of lots of fluid connections among large numbers of nodes. The structure is roughly modelled on the neurons and synapses of the human brain. As systems are 'trained' with images or other data, they adjust the strength of the connections within the neural network without any prior knowledge of the data being fed to them. After seeing enough examples, they automatically learn the characteristics of whatever it is they need to identify.

Predictive analytics

Statistical or machine learning techniques that analyse current and historical data to make predictions about future or unknown events. See also **Data Analytics**.

Recommendation engine

A computer system that tries to predict and rank user preferences, e.g. on movie, music, news and shopping sites. While these systems can be based on pre-programmed rules, they typically use statistical or machine learning techniques to discover similarities between users based on analysis of the data and choices of many users.

Resumé-screening software

A system, based on natural language processing and machine learning, that is able to read, evaluate, summarise, prioritise (and in some cases even devise appropriate interview questions for) job applicants, based on the content of their CVs and applications.

Self-driving cars

See autonomous vehicles.

Sentiment analysis

Software based on AI techniques such as natural language processing and machine learning, that analyses large amounts of unstructured data (such as social media feeds, blogs, etc.) to understand how people feel about a particular issue or product.

Smart robots

Robots, typically used in factories, that have camera or sensor technology and use AI techniques and technologies to perform a particular task, e.g. checking for defects.

Speech recognition (engine)

An AI-based system capable of understanding human speech. Thanks to **machine learning** and **deep learning** – and the rapidly growing amount of speech data available for analysis since the advent of voice assistants like Apple Siri, Amazon Alexa, Google Assistant and Microsoft Cortana – these are rapidly becoming far more accurate.

Strong AI

See **artificial intelligence**.

Turing Test

A test proposed by Alan Turing to ascertain computer intelligence. If a human could converse with a machine in natural language, on any subject, without realising it was not another human, then the machine would pass the test.

Virtual agent

See **chatbot**.

Virtual personal assistant

An AI chatbot designed to aid productivity by managing tasks such as organising meetings or setting reminders.

Weak AI

See artificial intelligence.

For more insights

Visit business.bt.com/solutions/insights

Offices worldwide

August 2020