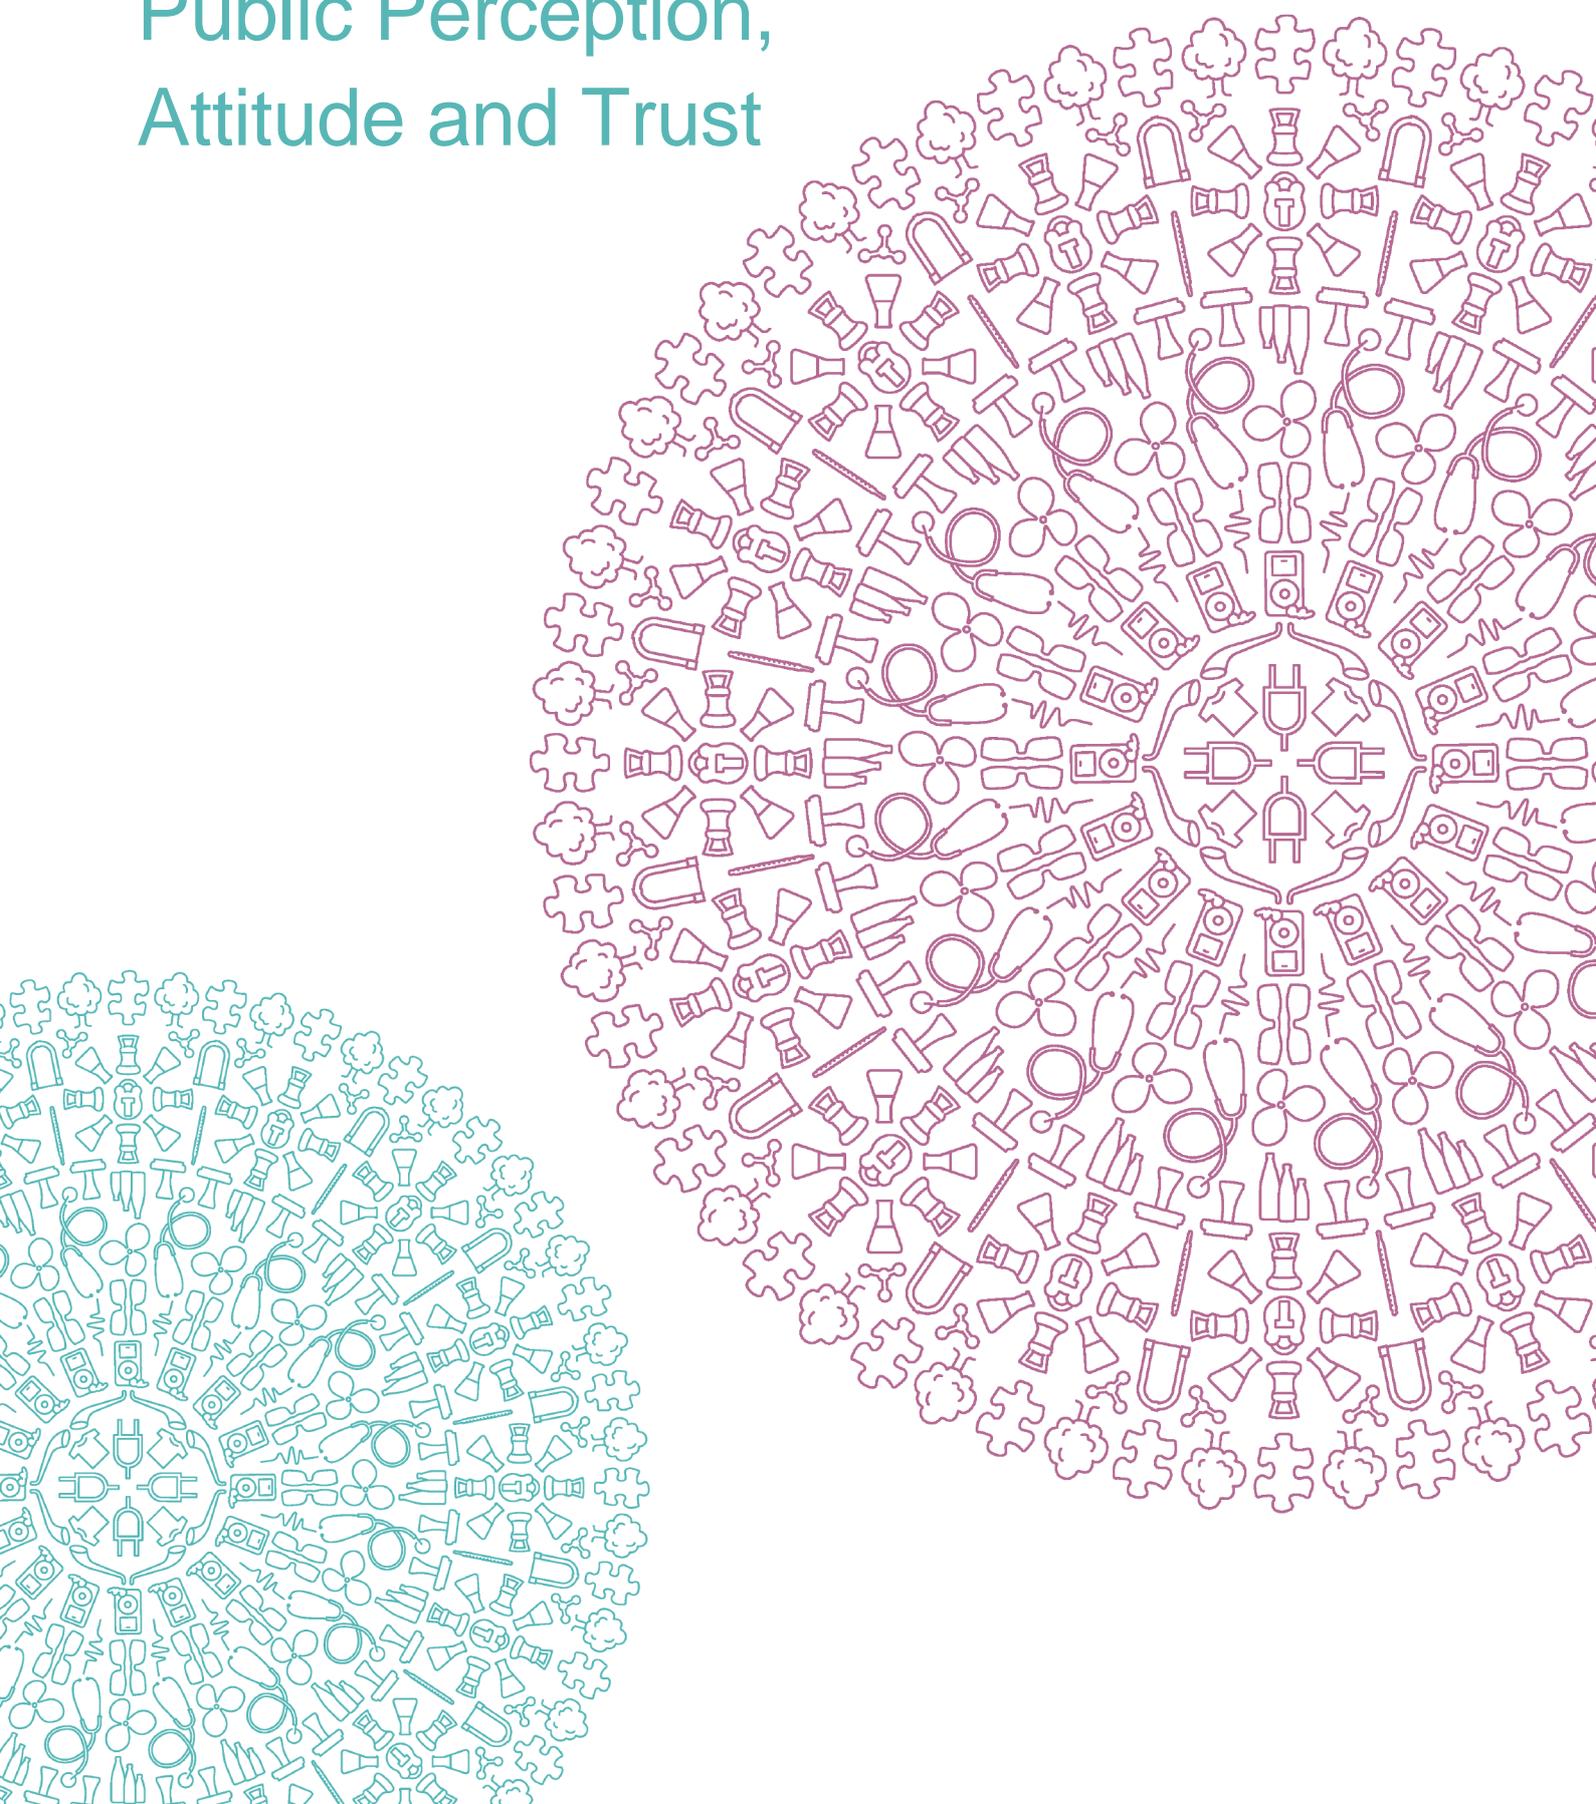


BRISTOWS

Artificial Intelligence: Public Perception, Attitude and Trust



Artificial Intelligence: Public Perception, Attitude and Trust

Introduction	2
Executive Summary	4
AI awareness and knowledge	6
Public perceptions of and attitudes to AI	9
Impact on jobs and the workplace	10
Use of personal data	15
Accountability, Responsibility and Regulation	17
Some other talking points....	18
Conclusion	20

Introduction

Artificial Intelligence (AI) is **out of the lab and into the real world**. It is changing the way companies do business, how governments provide public services, and how people engage with technology and each other. Businesses in all sectors are exploring how AI can help them to **reduce costs and improve services**, while the media is focusing its attention on **sensational headlines** predicting mass unemployment and other disasters.

What is clear is that there is an important debate to be had; one in which the AI industry (users, developers and other stakeholders) needs to play a public role. We at Bristows think this debate needs to put **public trust at its centre** – after all, widespread adoption of any new innovation can only happen if the people most affected by it accept its use.

In order to gain an insight into public understanding of AI, **Bristows has carried out a special survey** assessing people's perception of this technology, its potential uses, how it will affect their daily lives, and the role regulation might play.

We hope this survey and the insights set out in this whitepaper add to this debate and help the industry to focus on the importance of engendering public trust in this exciting, transformational technology.

The results certainly make for interesting reading.

About the survey

The survey was conducted by Censuwide on behalf of Bristows, with 2103 general consumers in the United Kingdom questioned between 12 July 2018 and 16 July 2018.

Censuwide abides by and employs members of the Market Research Society, which is based on the ESOMAR principles.

About the authors

Bristows' *AI and Robotics* team has been working in technology for over 25 years, with a focus on AI for the past 5. We have advised clients on various issues relating to the development and use of AI technologies, including drones, robotic process automation and healthcare data analytics. We have also been actively involved with UK Government and the European Commission, assisting in the development of ideas surrounding the implication of AI on law and regulation.

Our submission to the House of Lords Select Committee on AI published in April 2018 highlighted the importance of public trust to the successful propagation of AI systems across society, and that clarity of accountability, responsibility, and access to appropriate remedies when AI goes wrong will be vital to the success of the industry. You can find our full report [on the UK Parliament's website](#).

We continue to be at the forefront of developments in this exciting sector and to write about AI and everything technology-related at www.bristowcookiejar.com



Chris Holder

Partner, Head of AI and Robotics



Vikram Khurana

Senior Associate,
Commercial IT



Mark Watts

Partner, Head of
Commercial IT

CENSUSWIDE
THE SURVEY CONSULTANTS

About Bristows

Bristows is an independent, London-headquartered, European law firm that has advised innovative clients since its inception in 1837.

We help clients grow in life sciences, technology and other dynamic sectors; clients on the edge of tomorrow; those creating new technologies and ideas, and those embracing them.

Bristows offers an unusually deep knowledge of the industries it serves – we like to recruit inquisitive minds, many with science and technology backgrounds. We live and breathe our clients' business, can talk their same language and have a keen eye on the future.

Executive Summary

Artificial intelligence is not new; modern AI has its roots in Alan Turing's test of machine intelligence in 1950 and the term was coined by a Dartmouth College professor in 1956. Today, the term denotes a broad range of technologies, concepts and uses. In this whitepaper, we use the term AI to refer to a set of computer science techniques that enable systems to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making and language translation.

In the public consciousness, AI has been popularised in the form of science-fiction and Hollywood movies (often in the context of super-intelligent robots attempting to enslave the human race). This is changing fast with the advent of driverless cars and personal voice assistants such as Alexa and Siri bringing AI to the consumer market. Meanwhile, media outlets are feeding the hype and hysteria about the potential negative impact of the new technology, such as mass unemployment, restrictions on personal freedom, increased economic inequality and a devalued human experience.

The dichotomy between AI hype and AI reality means it is difficult to know what the general public really thinks about today's most disruptive technology. Through our survey, we have sought to gain a better understanding of the public perception of AI by asking more than 2,000 people from a variety of backgrounds and locations in the UK what their understanding, attitude and expectations of this emerging technology are. Our survey reveals a number of interesting insights that we think the AI industry as a whole (users, providers and other stakeholders alike) needs to understand when designing and implementing AI systems that affect the general public.

The following are the key findings from our survey:

- **Public understanding of AI is “broad” but not “deep”.**
Three-quarters of respondents either know what AI is, have limited knowledge of it or consider themselves experts. However, only 1 in 7 respondents believe they have had direct contact with AI, and only 2% think AI is already having an effect on society (which suggests it is not always clear to people when they encounter AI and how it is being used in the world around them).
- **Expectations are high, but certainly not all positive.**
While many respondents were able to identify certain AI abilities and functions available today, a significant number also think AI could perform tasks that are currently beyond the state-of-the-art. Meanwhile, 47.4% of respondents believe AI will have a negative effect on society.
- **Young people are most optimistic about AI.**
Respondents aged under 35 are more likely to believe they have had

contact with AI and to think it will have a net positive effect on society. This age group is also more likely to embrace automation in the workplace. Meanwhile, only one-quarter of respondents aged 55 or over were minded to automate the most repetitive part of their job.

- **Employment concerns exist, but potential workplace benefits are acknowledged.**

While a significant proportion of respondents would not consider using AI in their own job, many were also minded to do so if they could save time or reduce errors. Traditional blue-collar jobs were seen as most at risk of AI automation, while professions such as journalism, law and the creative industries were seen to be less likely to be affected.

- **Privacy and data protection implications are not well understood.**

Over half of respondents either thought AI would not use their personal data or did not know if it would or not. Half of respondents were not comfortable with their personal data being used by AI to perform tasks for them.

- **The AI industry should be accountable and responsible to the public.**

More than two-thirds of respondents believe AI should be regulated, with almost half looking to the UK government or supra-national regulatory bodies to take the lead in ensuring accountability. A significant proportion also want the AI industry as a whole to self-regulate in some way.

AI awareness and knowledge

The use of the term “AI” in society today is very broad. It is used to describe concepts as diverse as data mining to autonomous weapons systems. We consider robotic autonomous systems, big data, the Internet of Things, and AI to be different forms of technology under the umbrella of ‘Industry 4.0’.

This distinction may be clear to people operating in the AI sector, but it is not necessarily understood in wider society as our survey shows.

We asked our survey respondents to rate their personal knowledge of AI (Figure 1). The majority of respondents seem to have at least heard of AI, with almost 27% saying they know what is and almost 40% claiming to have “limited” knowledge. Only 9.1% of respondents said they had never heard of AI, while on the other end of the spectrum only 8.4% said they “know a lot about” or “were experts in” AI.

When asked about personal knowledge of AI

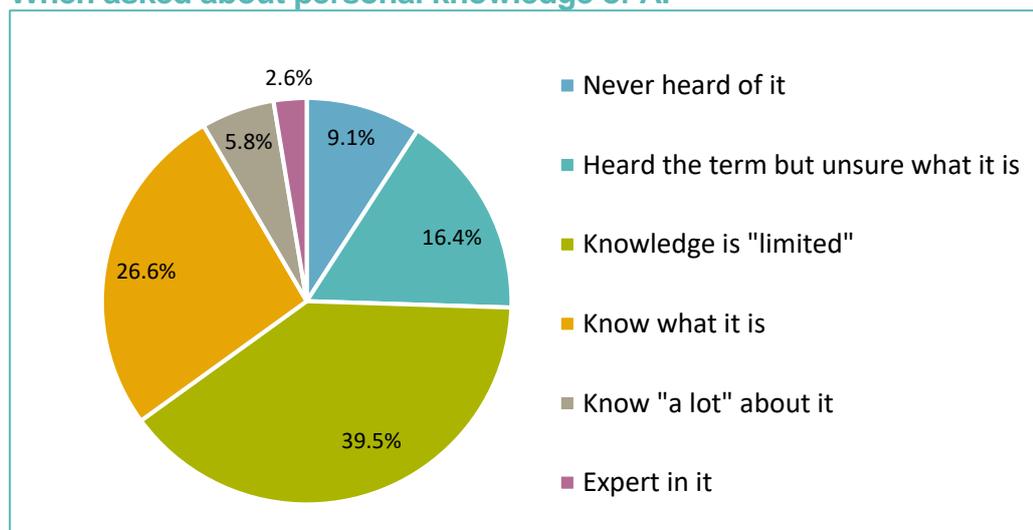


Figure 1

We knew that a self-assessment of one’s own knowledge of a topic as broad as AI would not necessarily be representative of the truth so we decided to test respondents’ knowledge of AI by asking them to select from a list the characteristics they thought best describe AI (Figure 2).

We included characteristics we consider to accurately reflect the level of AI available today (for example, AI does not necessarily need a physical form, it could just be software), and some that are taken straight from science fiction (for example, at this stage AI cannot self-replicate).

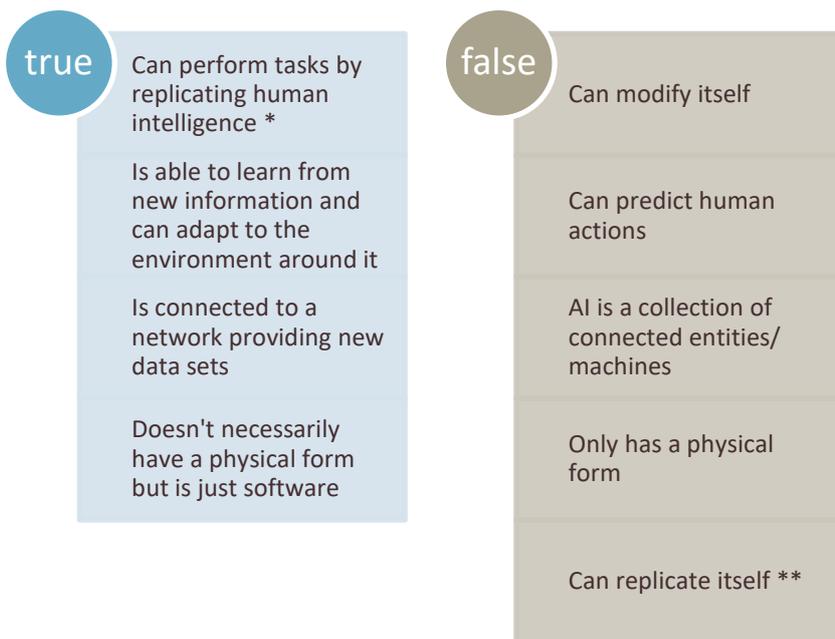


Figure 2

* Visual perception, speech recognition, decision making, translation between languages, generating text/speech that sounds natural.

** NB: A "none of the above" option was also available to respondents

Almost 1 in 5 (19.4%) of respondents thought that AI can modify itself, while 1 in 6 (17%) thought that AI can predict human actions.

These reactions were higher amongst respondents who said they "knew a lot about AI" (39% and 28% of them respectively also thought AI had the two characteristics mentioned above), which could suggest that available information about AI - through general and specialised media - is overestimating its current level of sophistication and therefore the type of application that AI is being currently used for/will be used for in the near-term.

We also wanted to see if people had a clear view of how widespread AI already is (Figure 3): do people realise that many financial products they buy are created by AI or that assistance chatbots they turn to ask simple questions are based on relatively straightforward decision-tree automation?

When we asked about their experiences with AI, over 3 in 5 (62%) respondents said that they had not been in contact with or used an AI application, while 23% did not know whether they had or not.

In comparison, only 1 in 7 (15%) of respondents said that they had been in contact with or had ever used an AI application.

This result suggests that AI can often take an 'invisible' or 'behind the scenes' form, potentially making it difficult for people to understand when they are interacting with AI in their daily lives.

Have you been in contact with or used an AI application?

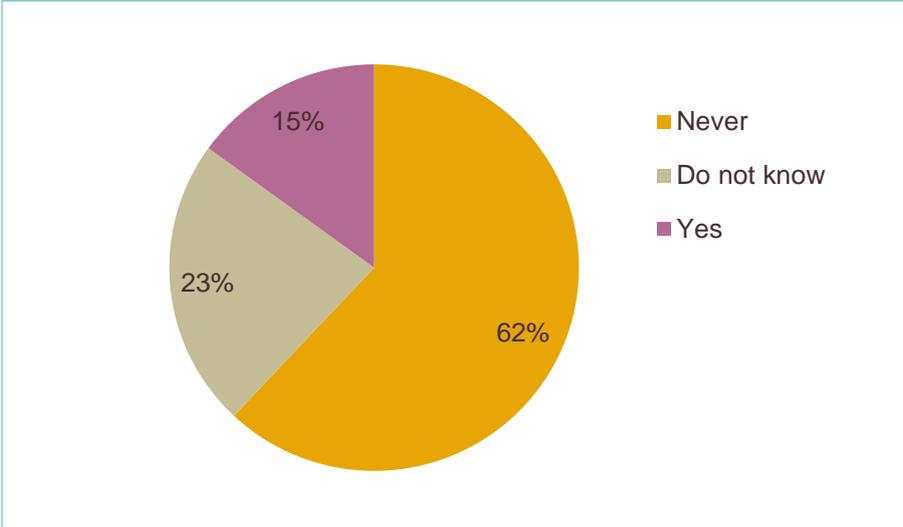


Figure 3

Public perceptions of and attitudes to AI

Having investigated what people think they know about the technology, our survey examined their general perception towards AI. We asked our respondents when they thought that AI would start having a positive or negative effect on society.

We provided the respondents with multiple options to choose from, from “It already has” to “Never”, and including “This year (2018)”, “Next year (2019)”, “Within the next 5 years”, “Within the next 10 years”, “Later than the next 10 years” or “Don’t know”.

The answers show that most people cannot gauge the speed at which AI is changing society (for the good or the bad), with more than one-third of the respondents saying they did not know when positive and negative effects might appear (34% and 45%, respectively).

On average, respondents thought AI would start having a positive effect on human kind in four years (4.32) and a negative effect in five years (4.74), with one in four (28.6%) saying positive effects would appear within the next five or between five and ten years and one in five (20%) saying negative effects would appear within the same time period. This indicates that people see AI as a “game-changing” technology and therefore expect to see results relatively soon; that this feeling is more marked in those who are worried about possible negative impacts suggests this effect is caused by adverse media reporting of the potential harm that can be done by AI.

Finally, perhaps underestimating the state of the market, or maybe because it is too early to attribute positive or negative effects to the introduction of AI in business and society, only around 4% of the respondents thought AI had already started having either a positive or negative effect on society (2.4% for positive and 2.2% for negative, respectively)¹.

We would suggest that increasing the public’s awareness of what AI and robotic autonomous systems are, and when and how they are likely to be used, will be helpful in managing the public’s expectations and reversing our inherent resistance to change.

The easiest way for the public to get used to and prepare for the advances in these new technologies is to see them in action, to use the applications and machines, to be educated about the benefits and not to be continually bombarded by stories of what may happen ‘when the robots take over’.

¹ A curiosity: while most geographic locations conform to the national average, people in Brighton have very different ideas: none of them think that AI has already started to have a positive effect on society, with 22.4% putting their money on the 2018-2022 period (within five years) instead.

Impact on jobs and the workplace

One of the most widely reported potential impacts of the use of AI has been on the world of employment and the way we work, and so we wanted to explore this more fully. Were people worried, or excited, about the impact of AI on their jobs?

Automation has been around for decades with workers used to using machines to help them do their jobs – from automated production lines to software spreadsheets and calculators – machines are used every day by the vast majority of people, to assist them in their daily tasks.

We asked respondents if they would use a solution to automate the most repetitive part of their job, and the answers are almost evenly split between people who thought yes (32%), people who thought no (38%) and people who did not know (30%).

When questioned regarding why they had given their positive answers, over 3 in 5 (62%) of respondents thought they would use it to save time, which would allow them to do more work, and/or focus on more fulfilling tasks instead.

The time-saving opportunity has the highest appeal to our respondents, and we think this is because it is also one of the most common, positive effects of the use of AI as set out in the media.

Why would you automate your job?

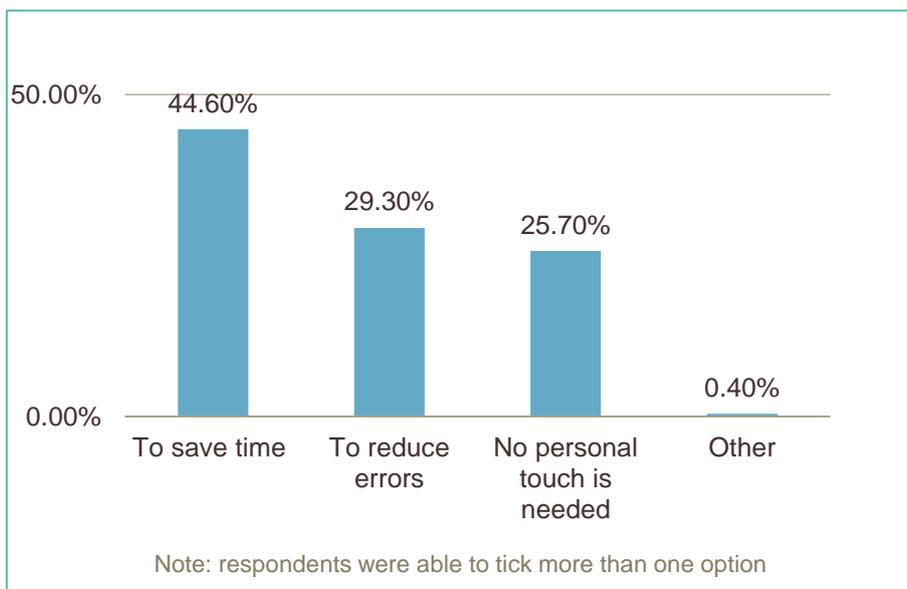


Figure 4

People who thought they would not want to use a potential solution to automate the most repetitive part of their job indicated the main reason for this was that

certain tasks required a personal touch or customisation and so would not benefit from the use of AI.

If you would not use AI as part of your job, why not?

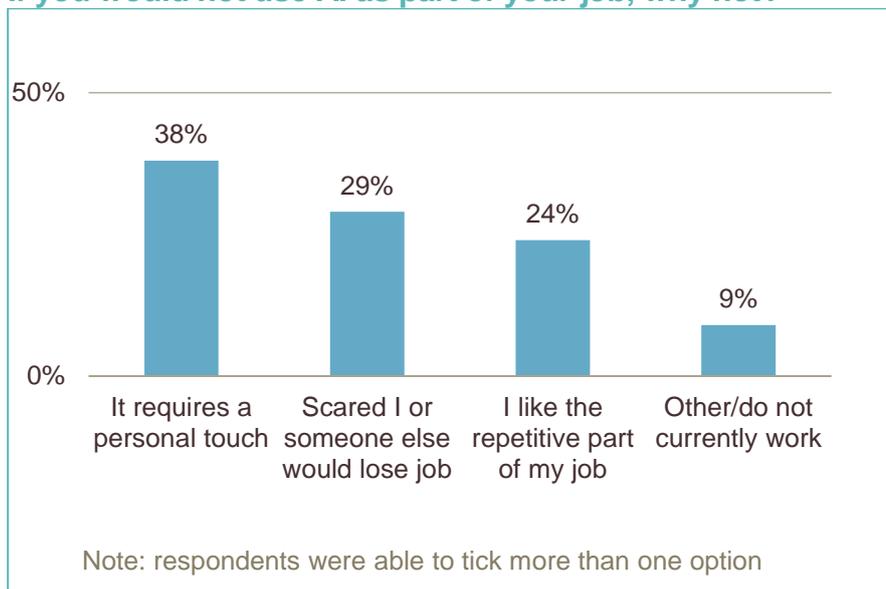


Figure 5

For many people, though, there will be no choice on whether they use AI or not. As companies start to introduce these new technologies across their businesses, the way their employees work will change.

What did our respondents think would be the jobs that would be most affected by AI? We presented our respondents with a list of some of the most common job roles in the UK² and asked them to pick the top three that would be most affected. Results show:

1. Factory workers (26%)
2. Cleaners (12%)
3. Office admins (12%)

We think these answers only partially reflect what the AI industry is working towards. Yes, a lot has been done to automate cleaning and logistic processes, but not many companies are focusing on office administration yet, while financial services companies are looking into it and investing heavily in artificial intelligence.

At the bottom of the list were those roles where our respondents potentially believed a human perspective was of more importance. The roles were:

² The full list included: accountant / tax professional, care worker, chef / cook/ kitchen staff, cleaner, creative professions (artist, designer), delivery person, doctor / vet, driver, factory worker, hotel manager, journalist / writer, lawyer, marketing/ public relations professional, nurse, office administrator, civil servant, transport worker, retail assistant, teacher. We included an option for 'other, please specify', so respondents could have included anything else they wished.

- Journalists (2.7%)
- Lawyers (3.4%)
- Chef / Cook/ Kitchen staff (3.7%)
- Care workers (3.9%)
- Creative professionals (artists, designers) (4.1%)

While there has been a lot of talk about the “Robo-Lawyers of the Future”, our colleagues will be glad to hear that the market is rather looking at a “Robo-Law-Assistant” that would help speed up research, assist with due diligence of documentation or draft documentation to support “paper-heavy” transactions, but it is still too early to suggest that AI will provide ‘legal advice’ without human involvement.

In what areas of work did our respondents believe that AI would be best suited?

Data analytics came top of this list, regardless of industry sector, with 62% of the respondents thinking it was likely that certain tasks would be automated (and 18% absolutely certain of it).

This was followed by:

- monitoring tasks (59% thinking it will be likely automated and 14% absolutely certain of it),
- agricultural support (58% thinking it will be likely automated and 11% absolutely certain of it) and then,
- household tasks for elderly or disabled people (58% with 12% absolutely certain).

On the other hand, only 34% of our respondents replied they could see emotional support services being provided by AI instead of humans. It is noteworthy that there is relatively strong push back on the idea that AI robots would provide sexual services with 49% of respondents stating they would not allow this.

Which industries do you think are likely to be automated?

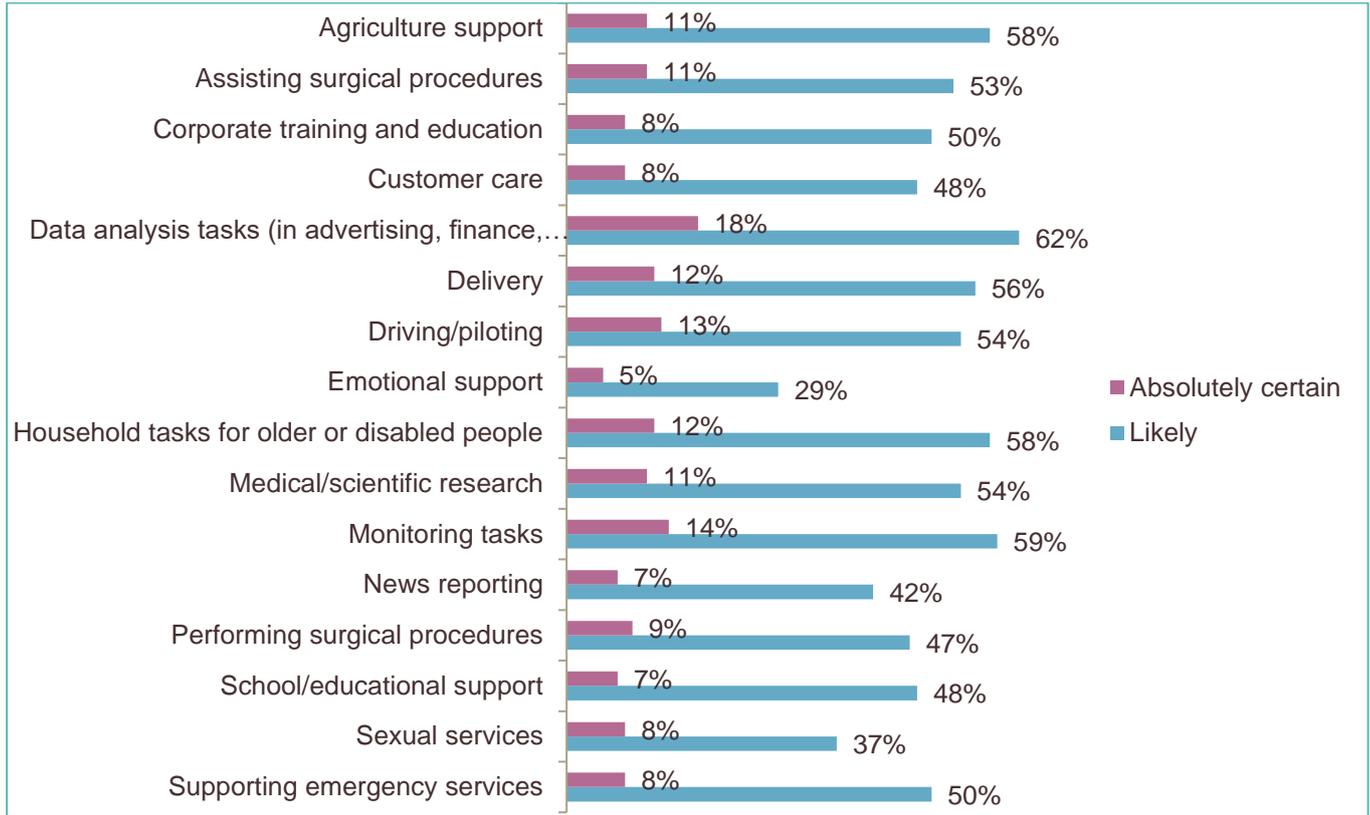


Figure 6

Most respondents (on average 44%) were not certain about using AI themselves for any of the tasks in the list of what they thought would be performed by AI in the future.

Those who could imagine AI being used for the provision of emotional support, sexual services, or driving tasks, were the least comfortable with using AI to perform these tasks for themselves.

Finally, people who thought that AI would be used for deliveries, monitoring tasks and household tasks for elderly or disabled people in the future were the most certain they would use the technology themselves.

From an industry investment perspective, perhaps it is these areas where more R&D and product development would be best focussed, as these appear to be areas favoured by our respondents for positive engagement and future use.

We think this is an important point: after all, an industry cannot thrive without customers. These answers show that roughly half of the public are potential users of future AI products, services and tools.

Would you allow AI to perform the following tasks for you?

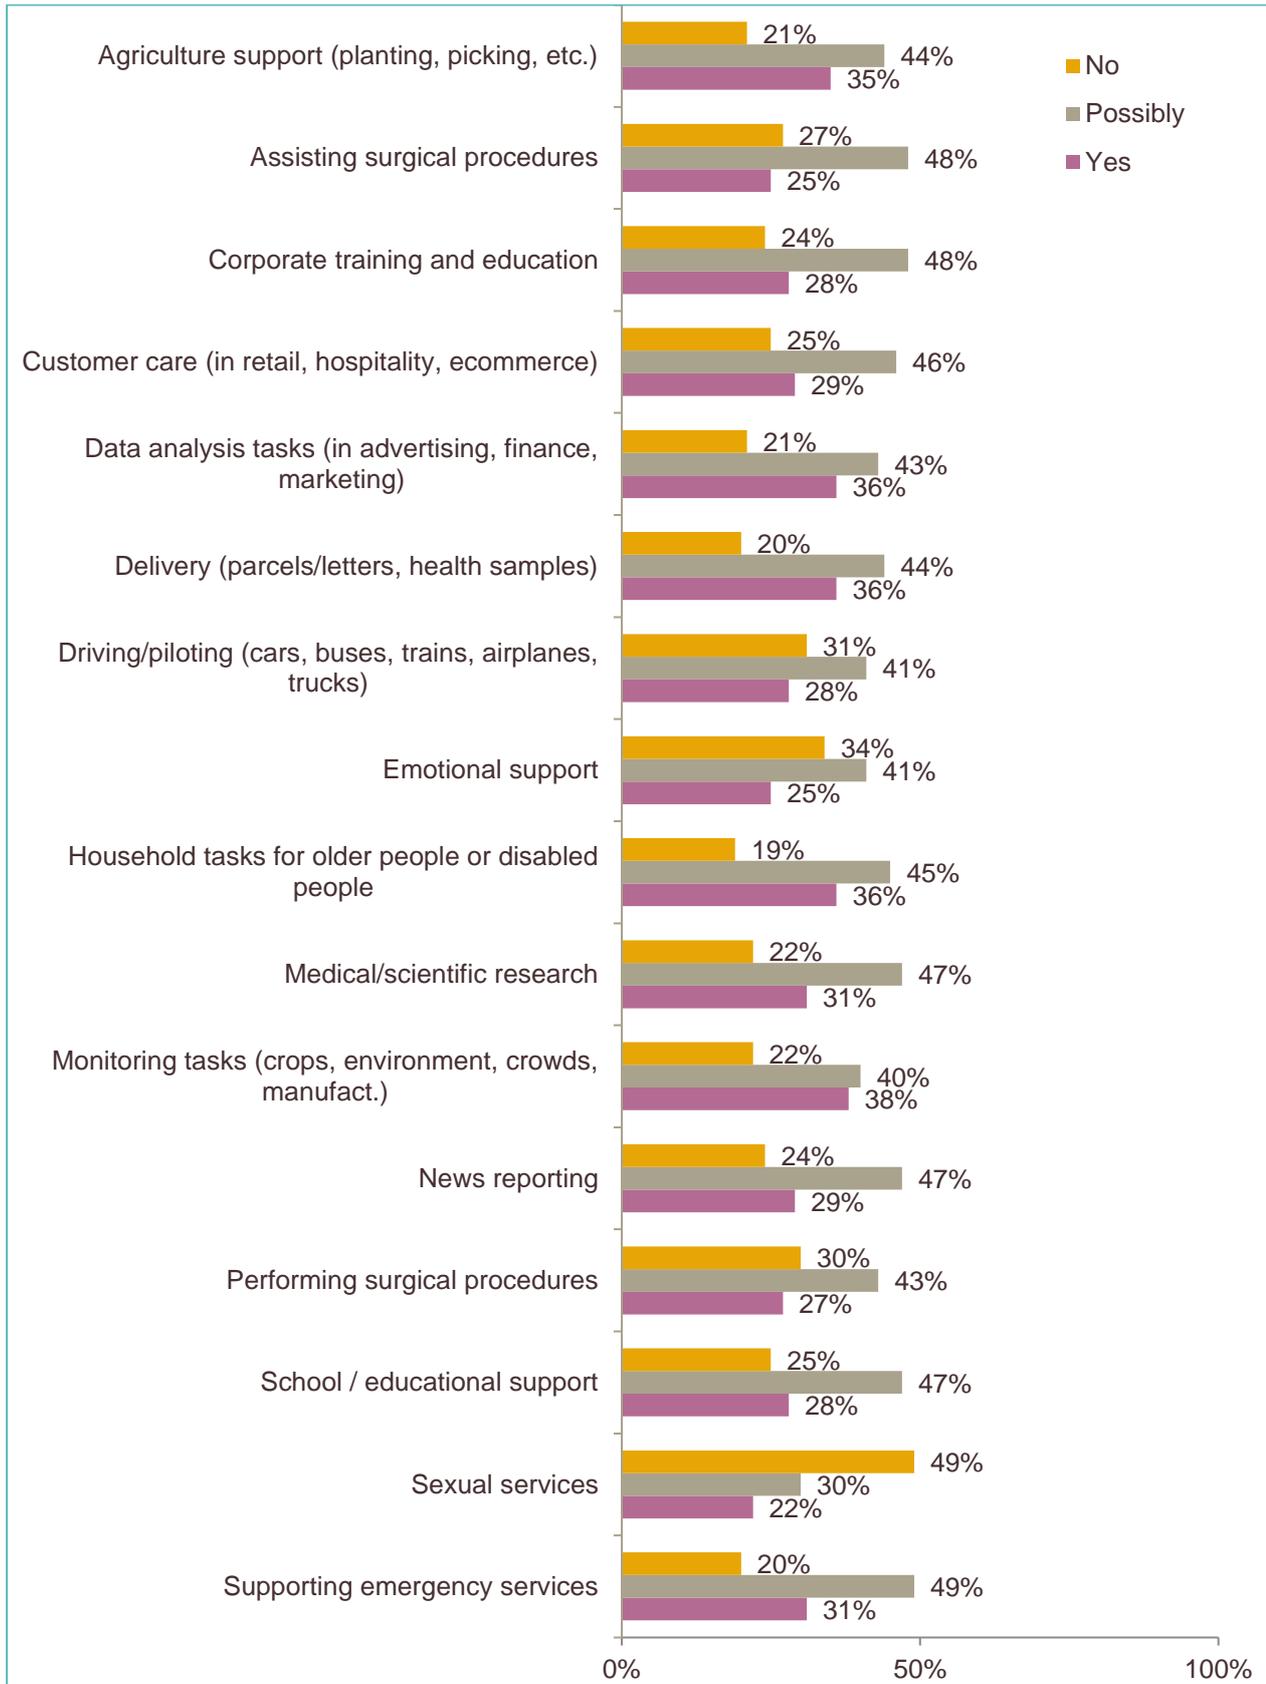


Figure 7

Use of personal data

AI systems rely on vast inputs of data in order to generate the outputs that businesses and people can use in their daily operations. Often this data will be personal data, perhaps relating to a person's age, gender, online shopping habits or route to work. For consumers, AI-powered apps applied to their personal data can be very useful or even essential in modern life: music apps can predict what artists you might like based on your listening habits; satellite navigation apps can guide you through the quickest and safest route to your destination. In other cases, companies can draw significant insights about their customer base or workforce from collecting and processing personal data through AI-based algorithms that assist their decision making.

We wanted to understand the extent to which the public is aware of and comfortable with the level to which their personal data is being used for these purposes. We already know people do not read the lengthy terms and conditions that govern their use of online services and generally do not alter privacy or security settings once they have signed up.

The responses to our survey showed that most people either did not think that personal data would be used by AI to perform tasks (22%) or did not know (39%) (Figure 8). Just under 2 in 5 (39%) of respondents thought that personal data would be used by AI to perform tasks for them.

Will personal data be used by AI to perform tasks?

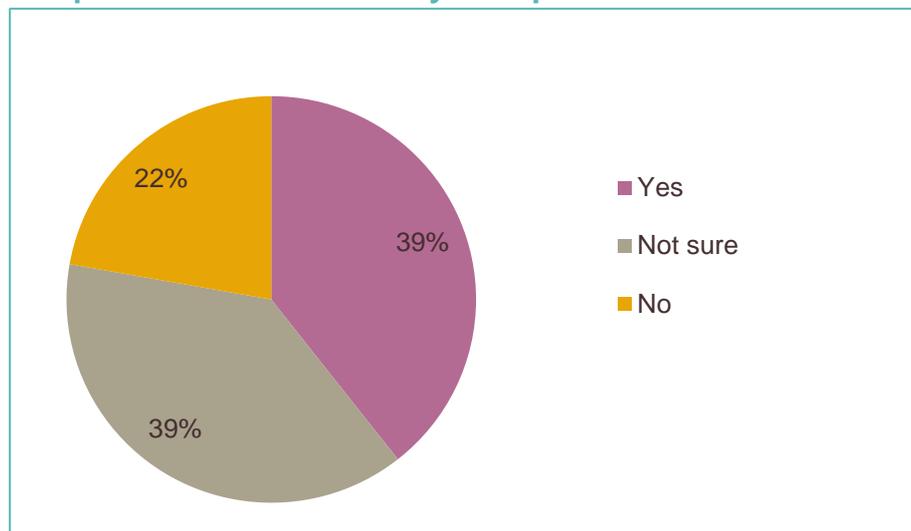


Figure 8

Just under a fifth (19%) of respondents were comfortable with the idea that their personal details and information regarding their personal life would be used by AI to perform tasks for them.

However, exactly half (50%) were not at all comfortable with their data being used by AI to perform tasks for them and this rose to 57% of respondents aged 55 and above, compared to 44% of respondents aged 16-24 (Figure 9).

Are you comfortable with your personal details and information about your personal life being used by AI to perform tasks for you?

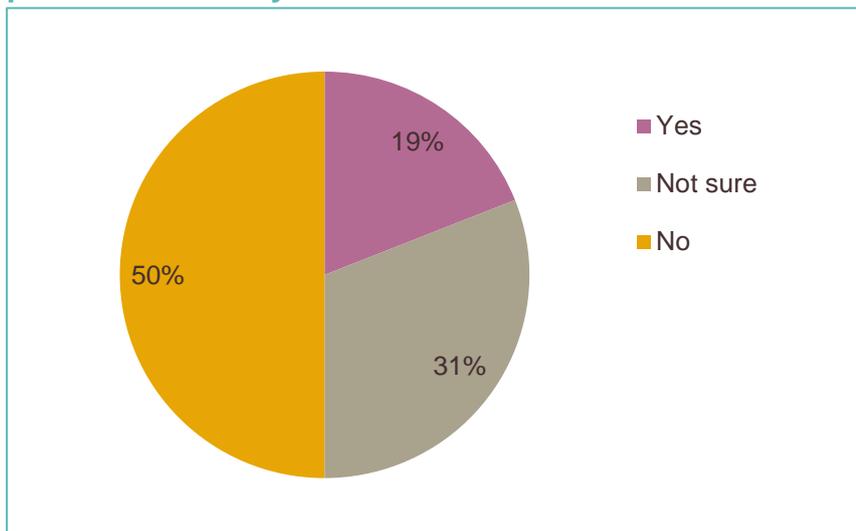


Figure 9

The results suggest a lack of awareness of when and how people's personal information is collected and processed by the services on which many of them have come to rely. When presented with the reality, many people express discomfort with this. This suggests the AI industry needs to be clearer with people about how their products and services work, and to make the case for the benefits to the individual and, in some cases, for example healthcare analytics, society at large.

Accountability, Responsibility and Regulation

There is a view that industries that are heavily regulated are more likely to be trusted by the public than others where there is very little or no regulation at all. For example, the airline industry is heavily regulated, with safety being of paramount importance – this is the basis on which people are willing to travel by air. There are parallels with innovative technologies like AI with real or perceived safety issues – it is possible that imposing regulation in whatever form could in fact engender the level of public trust that will lead to increased use and acceptance of AI among the public.

Our survey appears to bear out this hypothesis: 69% of respondents believed that AI should be regulated. While our view is that this reaction should not necessarily lead to a rush to regulate, it may be that in specific circumstances (e.g. sector-by-sector) some level of regulation may persuade the public that the industry is treating safety issues seriously (even if it is already doing so in reality).

Of those calling for regulation, a large proportion of respondents (42%) thought this should be a task for UK central government. 15% also thought the EU or UN should be responsible – this relatively low proportion seems to reflect a post-Brexit trend away from devolving national law-making to supra-national bodies.

Finally, almost 1 in 5 respondents thought the AI industry itself should self-regulate in some way – this lends itself to the type of industry standards and codes of practice prevalent in established industries that have managed to avoid regulation stifling their development.

Who should be responsible for regulating AI?

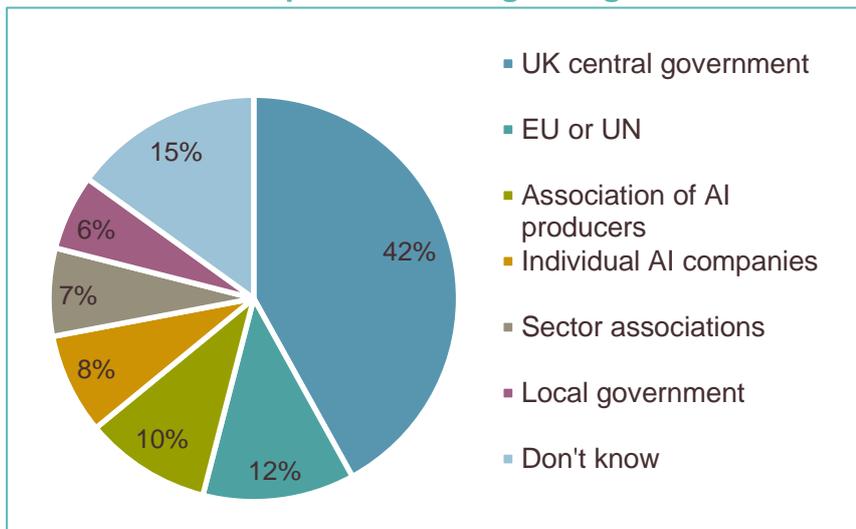


Figure 10

Some other talking points....

Age has been shown to be a differentiating factor in many of the answers that we received, with different age groups having differing views. We have already seen in a previous paragraph that younger people are more comfortable with sharing their data with AI applications, but there are some other interesting areas.

First of all, most people appeared unsure on what AI means, but older people had a different level of awareness to those of the younger generation:

- Just under half (49%) of respondents aged 55 and above said that their knowledge of AI was limited;
- 2 in 5 (40%) of respondents aged 45-54 said that their knowledge of artificial intelligence was limited;
- 30% of respondents aged 16-24 said that their knowledge of AI was limited.

Out of each age group, how many said their knowledge is limited?

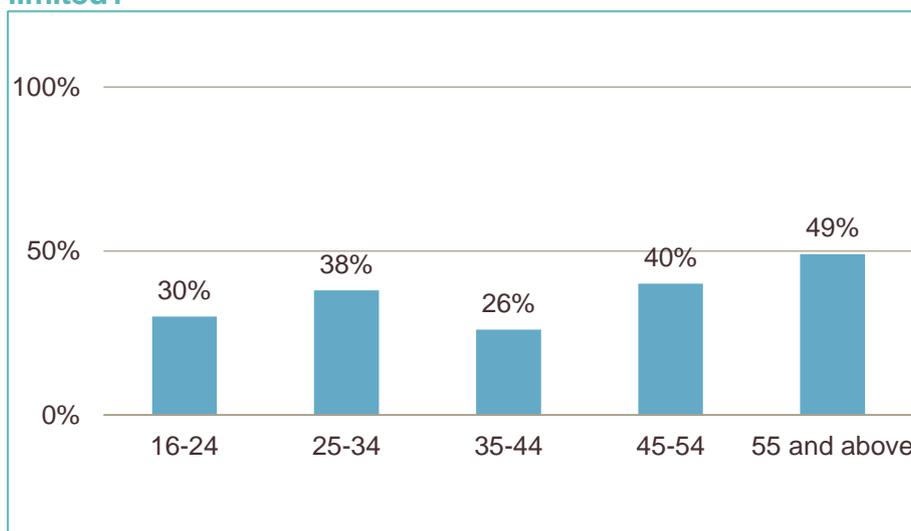


Figure 11

Secondly, the number of people who said that they had not been in contact with or used AI was over 7 in 10 respondents (74%) aged 55 and above, while younger people in the 16-24 and 25-34 age brackets had a higher percentage of positive engagements with AI.

Age split for people who thought they had not been in contact with AI

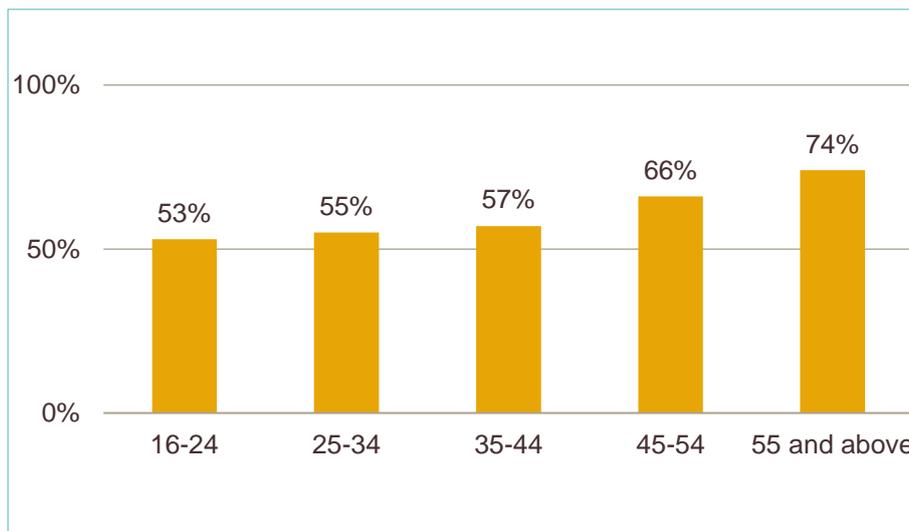


Figure 12

Finally, with regards to the issues surrounding the likely impact of AI, different age groups showed different views.

Respondents aged 55 and over thought that AI would start having a positive effect on society in just over 5.5 years (5.51) on average, while those respondents aged 16 – 24 thought that such positive effects would be achieved in just over 3 years (3.4) on average.

Conclusion

How much misunderstanding about AI, then, is there among the general public? And to what degree do people's hopes and expectations – for better or worse – match up to reality? Many of our survey respondents indicated an awareness of AI technology but at the same time seemed unaware of how and where they encounter it in their daily lives. Younger respondents tended to embrace AI in their lives more than older respondents, but many from across the age groups retain concerns about privacy, accountability and the risk to their jobs. Also, while it is clear that people think AI will ultimately have an impact on their lives, they differ in whether the effects will be overall positive or negative and the time period over which change will start to appear.

The range of views and opinions our survey reveals among the public appears to reflect the divergence that exists between the media hype cycle and the views of experts in the AI field – indicating that (as with many other issues in our society) what the public thinks about AI is at least partially influenced by the messages they receive from reporting on this area. That means the AI industry as a whole (including the companies developing the technology, and the organisations putting AI at the forefront of their agenda) will need to increase public confidence in order to obtain the wider buy-in needed for the successful development of AI in society.

We think the results of our survey point to several ways in which the AI industry might go about increasing public trust and confidence in AI systems that affect people's lives:

- Firstly and perhaps foremost, the public is confused about what AI is, and where and how it is used in the world around them. To address this, the industry should be more **open about embedding AI in products and services**, and to be **upfront about how AI works** in a way that recognises what is currently possible and what promises remain out of reach.
- Secondly, concerns about negative AI impact on jobs are prevalent. While it is clear that AI will have an **effect on people's jobs** now and in the future, employers could do more to foster the use of AI in a way that does not necessarily equate to job displacement, but rather **augments what individuals do in their daily tasks** to free them up to perform more creative and value-contributing tasks.
- Thirdly, people do not necessarily understand **the way AI will use their personal data** – and when they do, they may not be comfortable to share it. Other recent research has revealed that people are generally happy to share their data if they believe that there will be societal benefits (e.g. with the NHS and its partners in order to provide better treatment and health outcomes), so the challenge is to convince people that AI is something not to be feared but rather it is something to embrace and work with. The

industry should be clear that access to personal data is required in order for some (but not all) AI systems to deliver public benefits, but that its use is **consistent with data protection laws** and (where this is accurate) **embeds key privacy principles** such as data minimisation and data portability.

- Finally, many people are looking to **outside bodies** to step in to **regulate AI**. The AI industry should develop its own internal codes, practices and kite-marks to provide comfort to people that **AI systems will be safe**, and to allow them to distinguish AI that should be safe to use from AI that has not passed regulatory muster. This approach could also limit the type of over-regulation that can stifle innovation in a nascent industry.

We at Bristows are AI advocates and we believe artificial intelligence has the potential to improve how companies do business, how governments provide public services, and how people engage with technology and each other. We hope the industry, and all relevant stakeholders in the AI race, take the steps required to **promote public confidence and trust** that will be required in order to ensure the potential of this transformative set of technologies are fully realised.

BRISTOWS

This material is for general information only and is not intended to provide legal advice. For further information, please speak to your usual Bristows contact or visit our website: www.bristows.com

Bristows LLP
100 Victoria Embankment
London EC4Y 0DH
T +44(0)20 7400 8000

www.bristows.com

