



Department  
for Transport

**Using AI in consultations and correspondence**

**Thinks Insight & Strategy research report**

September 2023

Department for Transport  
Great Minster House  
33 Horseferry Road  
London SW1P 4DR



© Crown copyright 2023

This publication is licensed under the terms of the Open Government Licence v3.0 except where otherwise stated. To view this licence, visit <https://www.nationalarchives.gov.uk/doc/open-government-licence/version/3/> or contact, The National Archives at [www.nationalarchives.gov.uk/contact-us](http://www.nationalarchives.gov.uk/contact-us).

Where we have identified any third-party copyright information you will need to obtain permission from the copyright holders concerned.

This publication is also available on our website at [www.gov.uk/government/organisations/department-for-transport](http://www.gov.uk/government/organisations/department-for-transport)

Any enquiries regarding this publication should be sent to us at [www.gov.uk/government/organisations/department-for-transport](http://www.gov.uk/government/organisations/department-for-transport)

# Contents

About and acknowledgments	5
About Thinks Insight & Strategy	5
Acknowledgements	5
Executive summary	6
Background and summary of methodology	6
Key findings	6
Key insights to take away from this research	7
1. Background	9
1.1 Background	9
1.2 Research objectives	10
1.3 Research methods	10
1.4. Sample	12
2. Spontaneous views of AI	16
3. Overarching views of the use cases	21
3.1 Comfort with the use cases	21
3.2 Factors affecting views of the use cases	22
3.3 Stakeholder views	25
4. Perspectives on Use Case 1: Using AI to analyse consultation responses	27
4.1 Public response to use case 1	27
4.2. Stakeholder response to Use Case 1	30
5. Perspectives on Use Case 2: Using AI to respond to correspondence	32
5.1 Public response to Use Case 2	32
5.2 Stakeholder response to Use Case 2	36
6. Guiding principles of the application of AI	38



# About and acknowledgments

## About Thinks Insight & Strategy

Thinks Insight & Strategy is an international insight consultancy focused on providing our clients with the insight they need to make better decisions. We do this by putting the people that matter most to our clients at the heart of their thinking.

[www.thinksinsight.com](http://www.thinksinsight.com)

## Acknowledgements

We would like to extend special thanks to:

- Dr Jonathan Bright, Fellow at The Turing Institute
- Ben Lyons, Head of External Affairs and Innovation at the Centre for Data, Ethics and Innovation

# Executive summary

## Background and summary of methodology

The Department for Transport (DfT) is exploring how it can apply artificial intelligence (AI) and machine learning (ML) to improve efficiency and quality across its work. In particular, DfT has been exploring how generative AI could be used in two different scenarios: 1) the analysis of consultation responses and 2) supporting the drafting of responses to external correspondence.

Consequently, DfT commissioned Thinks Insight & Strategy to conduct research with the public to explore:

- 1.1 Public knowledge around AI and its various applications.
- 1.2 Sentiments around the application of AI with regards to the specific use cases.
- 1.3 Public opinion regarding the benefits and risks involved with the application of AI.
- 1.4 The best ways to communicate with the public about these use cases.

To achieve these aims Thinks Insight & Strategy undertook a multi-stage research approach including: three deliberative workshops in London, Newport and Glasgow (each with 24 members of the public), 10 in depth interviews with engaged citizens and nine in-depth interviews with DfT stakeholders.

## Key findings

This research revealed eight key findings:

- 1.1 Despite low confidence in their knowledge of AI, participants believed that AI will be an important technology of the future. They therefore believe it is inevitable and appropriate that the UK Government (and consequently DfT) will utilise it for a variety of tasks.

- 1.2 Participants spontaneously believed that the key benefits of AI use are speed and efficiency and the key risks associated with its use are concerns around job losses, humans becoming overly dependent upon it and a loss of a 'human touch'.
- 1.3 Both use cases (consultations and correspondence) were low salience activities for the general public. Whilst these participants were able to articulate views about each use case, in the context of low salience it is unlikely that either proposed use case will elicit strong reactions from the general public outside of a research context.
- 1.4 Views of the benefits of AI remained consistent when discussing both use cases. These benefits were most compelling when focused on how they will impact the public rather than DfT (e.g. a faster response to correspondence is more compelling than the task taking up less of DfT's time).
- 1.5 Views of the most important risks shifted somewhat when discussing specific use cases with participants focusing much more on quality and accuracy of outputs. There was a consensus that increased speed was only valuable if AI was producing high quality outputs.
- 1.6 When they were first shown the two use cases, participants tended to be much more comfortable with AI being used internally at the Department to analyse consultation responses rather than for being used to draft external correspondence. This was driven primarily by a concern about AI's ability to draft empathetic responses to more emotional messages.
- 1.7 DfT's proposed mitigations addressed some of the participants' concerns, meaning that most ended the sessions feeling comfortable with AI being used for consultations. However, some participants wanted more reassurance on the topic of correspondence including whether humans would properly quality check AI outputs.
- 1.8 In line with their responses to the use cases, participants' guiding principles for DfT use of AI focus on quality, accountability and transparency.

## Key insights to take away from this research

1. There is wide recognition that AI is an important technology of the future and that DfT will utilise it.
2. The public are receptive to AI's use in saving time and driving efficiencies, although there is recognition that the benefits of speed can be lost if there isn't a focus on accuracy. In this context, the public highlighted the need for monitoring the quality of outputs.
3. The public are generally happier with the use of AI in areas not traditionally expected to need empathy, for example data analysis.
4. There is broad public acceptance for both use cases tested contingent on a focus on quality. However, some participants remain highly sceptical particularly using AI to draft external correspondence which is perceived to require more empathy.

5. The results highlight it is likely the public will hold an AI system to a higher level of accountability than humans.
6. When offered various options for how the use of AI would be communicated to the public, they supported the highest degree of transparency, for example: "To improve efficiency, this response was drafted with assistance from AI. It has been fully checked by a human".

# 1. Background

## 1.1 Background

The Department for Transport (DfT) is exploring how it can apply Artificial Intelligence (AI) and Machine Learning (ML) to improve efficiency and quality across their work. The DfT is interested in establishing itself as a leading department in the application of Artificial Intelligence (AI). To this end, DfT commissioned Thinks Insight & Strategy to develop a research project that would allow it to gain insight into the public's perceptions of the application of Artificial Intelligence and Machine Learning across its work. DfT identified two (2) use cases where AI could likely improve the efficiency and quality of the department's work, and would be used in direct contact with the public:

(a) Analysing consultation responses for internal use within DfT: 'consultations' provide an opportunity for members of the public, businesses and other organisations to share their comments on plans or ideas which are being considered by UK Government to help inform decision-making.

(b) Supporting the drafting of correspondence: letters and emails from members of the public, businesses or other organisations which require a response from officials.

Both use cases involve the analysis of large volumes of 'free text', defined as "text created by users who are free to type in whatever they want, rather than following some specific structure or limited set of responses". Both consultation responses and correspondence result in thousands of pieces of writing which currently need to be manually reviewed by members of staff. These circumstances can create issues around accuracy and 'human error', which are to be expected when staff are reviewing such vast quantities of free text. For example, staff may overlook or omit key information and insights as humans have a limited capacity for processing such large volumes of data.

DfT believe that the introduction of AI could significantly decrease the amount of resources and time that it currently allocates to the interpretation of free text. This would allow the Department to streamline both the consultation and correspondence processes and ultimately help the Department reallocate staff time and other resources in a more efficient manner.

Given that DfT serves the public, and that the public are active participants in the submission of correspondence and consultation responses, DfT understands it is of great

importance to gather and take into consideration their views of such proposed changes. As such, this research project is the first stage in exploring the public's perceptions of the application of AI across DfT's work, in the hopes that they can feed into the Department's plans to move forward with this technology in a way that best addresses the Department's and the public's needs.

## 1.2 Research objectives

This research aimed to capture the views of both members of the public and DfT's stakeholders on the use of AI to analyse consultation responses and respond to correspondence. The core research objectives were:

- (i) To assess the public's understanding of AI and its various applications both within everyday life and within UK Government departments
- (ii) To explore public sentiment regarding the application of AI in relation to this project's specific use cases
- (iii) To capture the public's assessments of the benefits and risks associated with the application of AI
- (iv) To determine the best ways to communicate with the public about these specific use cases

## 1.3 Research methods

This project utilised qualitative research methods in order to explore public and stakeholder perceptions around the core use cases. The research instruments, including all stimulus which was later shown to participants, were reviewed prior to use by two independent experts to ensure it was accurate and publicly accessible. The independent experts were: Dr. Jonathan Bright, Fellow at The Turing Institute and Ben Lyons, Head of External Affairs and Innovation at the Centre for Data, Ethics and Innovation.

The research engaged 91 participants across three strands of data collection:

- Three 3-hour deliberative workshops with 71 participants across the UK
- 10, 60-minute in-depth interviews with 'engaged' citizens (i.e. members of the public who have either responded to a UK Government consultation or sent correspondence to a UK Government department their local council in the past 2-years).
- Nine 60-minute in-depth interviews with DfT stakeholders.

All fieldwork was conducted between 13th July - 10th August 2023.

### The Deliberative Workshops

Deliberative workshops are designed to allow participants to explore an issue in depth through a process of information sharing, discussion, reflection and finally the presentation of their considered opinions. As such, they ensure that the public is able to input into

decision-making processes in a real and impactful way. This research included three deliberative workshops — one each in the cities of Glasgow, London, and Newport. Each workshop lasted three hours and had 24 participants who were working within smaller groups of six, so as to allow everyone the space to be able to actively contribute and share their opinions. The participants were carefully selected so that the sample reflected the demographic population of the UK in terms of age, gender, ethnicity, and financial background. Additionally, participants were also screened to ensure that there was a mix of familiarity with the concept of AI so that the workshops included persons who ranged from early adopters of AI technology to those who had limited exposure and knowledge of it.

The deliberative process commenced prior to the workshops, as participants were asked to complete an online pre-task. This short 20-minute exercise assessed their baseline understanding of AI and their views on the potential benefits and concerns around the use of AI technology overall. Participants were asked to share their initial level of comfort with the notion of the UK Government using AI, as well as their ideas of the best ways that they could integrate AI into their work. Some of this pre-task is quantifiable and has been aggregated and presented in section 2.

The flow of the workshop itself included:

- Reviewing the pre-task exercise and inviting participants to share further thoughts and their spontaneous views of AI.
- An information session on AI, exploring its current uses in daily life, its current capabilities and a live demonstration of Chat GPT.
- Reviewing the first use case, analysing consultation responses. Participants were introduced to the nature of consultations and both the volume and types of consultation responses the Department receives. The moderator then explained both the current process for analysing consultations and the more streamlined approach that would be implemented if AI were introduced.
- Participants then discussed the benefits and drawbacks of the use case in breakout discussions both spontaneously and then prompted. This discussion also included mitigations DfT offered to alleviate potential risks.
- The process was then repeated for the second use case, responding to correspondence.
- Following discussion of both use cases, participants were asked to reflect on everything they had discussed that evening and worked together to develop guiding principles for DfT when rolling out AI.

As part of the workshop participants recorded their sentiment via completing worksheets at various stages. This included at the beginning and at the end of the workshop. We were then able to understand how sentiment shifted during the course of the evening.

## Citizen Interviews

In addition to the workshops, 10x 60-minute interviews were conducted with 'engaged' members of the public. To be eligible to participate, interviewees needed to have either responded to a consultation or submitted correspondence to a local or central UK Government department. It was important to include their views in the research to understand how having first-hand experiences of the use cases influenced their views of

use of AI. The interviews began by exploring the interviewees' personal experience with either correspondence or consultations. They were then introduced to the relevant use case and walked through the both the current processes and the proposed plans following the introduction of AI. The interviewee was then prompted to share their perspectives on benefits, drawbacks, and mitigations, alongside their guiding principles for strategies DfT ought to adopt when implementing the technology.

## **Stakeholder Interviews**

This research also included nine x 60-minute interviews with stakeholders. These stakeholders were recruited with the assistance of DfT, who reached out to senior staff (e.g. CEOs or Directors of Strategy) within key organisations which engage with the Department on a regular basis. These interviews began by inviting stakeholders to share some information about their organisation's own experience integrating AI technology into their work and the benefits and challenges associated with that. They were then asked to detail their organisation's processes for both scripting formal correspondence and responding to consultations. Much like the citizen interviews, the sessions then explored the current process surrounding the use cases, as well as the plans following the application of AI. The stakeholders were then invited to share their opinions, concerns, and proposed guiding principles surrounding both use cases.

## 1.4. Sample

### Deliberative workshop sample:

<i>Demographic category</i>	<i>Sub-group</i>	<b>Number of participants</b>	
	Total	71	
<i>Location</i>	London	24	
	Newport	24	
	Glasgow	23	
<i>Gender</i>	Male	39	
	Female	32	
<i>Age</i>	18-20	2	
	21-30	13	
	31-40	20	
	41-50	10	
	51-60	14	
	61-70	10	
	71-80	2	
	<i>Ethnicity</i>	White	49
		Ethnic minority	22
<i>Socio-economic grade</i>	AB	19	
	C1	19	
	C2	18	
	DE	15	

**Citizen Interview sample**

<i>Demographic category</i>	<i>Sub-group</i>	<b>Number of participants</b>
	Total	10
<i>Location</i>	England	5
	Wales	1
	Scotland	4
<i>Gender</i>	Male	4
	Female	6
<i>Age</i>	21-30	3
	31-40	2
	41-50	2
	51-60	2
	61-70	1
<i>Ethnicity</i>	White	7
	Ethnic minority	3
<i>Socio-economic grade</i>	AB	5
	C1	4
	C2	1
<i>Type of engagement</i>	Responded to consultation	5
	Submitted correspondence	5

## Stakeholder interview sample

<i>Demographic category</i>	<b>Sub-group</b>	<b>Number of participants</b>	
<i>Location</i>	Total	9	
	England	9	
	Wales		
	Scotland		
<i>Sector</i>	General Transport	2	
	Rail services	1	
	Automobiles/motor manufacturing	2	
	Bus services	2	
	Aviation	2	
	<i>Role</i>	CEO/Director	2
		Chief information Officer	1
		Head of Innovation	3
Board Trustee		2	
Policy Adviser		1	
<i>Type of engagement</i>	Responded to consultations	9	
	Submitted correspondence	9	

## 2. Spontaneous views of AI

### 2.1 Knowledge of AI

All participants had heard of the term ‘artificial intelligence’ prior to the research. However, very few were confident in their own understanding of it. As shown below, when asked to complete a self-assessment prior to the workshop, none of the participants felt they understood the topic ‘very well’. Reflecting this low level of knowledge, participants across the deliberative sessions expressed shock and surprise at the information about the ways in which AI is already being used across various technologies that they were more familiar with (e.g. Amazon Alexa, Google Maps, etc.).

On a scale of 1 to 5 where 1 is not well at all and 5 is very well, how well if at all do you feel you understand Artificial Intelligence (AI) technology?

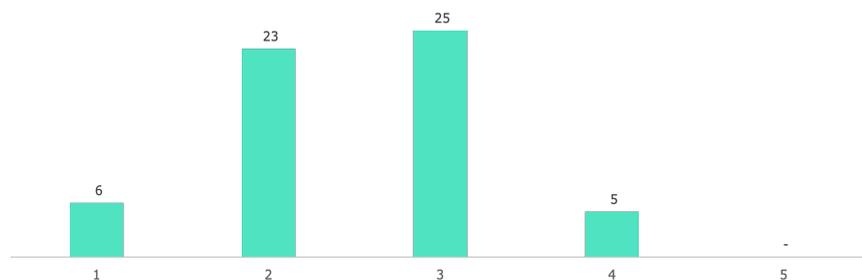


Figure 1 Chart showing responses to the question shared via an online survey prior to the workshop: "On a scale of 1 to 5 where 1 is not well at all and 5 is very well, how well if at all do you feel you understand artificial intelligence (AI) technology?". Base: 64

These findings reflect other research into public knowledge of AI. In recent polling on AI by Thinks Insight & Strategy whilst 66% of the public said they had heard of AI chatbots (like Chat-GPT, LLaMA) only 7% said they were regular users of this technology.<sup>1</sup> This polling shows a number of demographic groups appear to be more likely than the general public overall to claim to have heard of AI chatbots, including men (73% aware), 18-44s (75% aware, rising to 85% of 18-24s), ABC1s (75% aware), Londoners (77% aware) and people from ethnic minority backgrounds (83% aware).<sup>2</sup>

<sup>1</sup> Thinks Insight & Strategy, polling on public attitudes about AI (fieldwork conducted 23rd to 25th June 2023).

<sup>2</sup> Ibid.



Across all strands of research participants pulled out speed and efficiency as the key potential benefits of AI technology. There was a widespread belief that AI will be able to deliver work faster and more efficiently than humans are able to. This is reflected in wider research from The Turing Institute, where the public identified increased speed and efficiency as the top benefit of AI technology<sup>3</sup>.

---

*"[A benefit is] time efficiency of using AI to do tasks quickly for you...[It's] convenient to use and apply as it can process large amounts of data with minimal costs." - Workshop participant, Newport*

---

A smaller number of participants also highlighted that AI is likely to be able to deliver work of a higher quality than humans as it will eliminate human error (for example, as a result of getting tired or bored whilst working).

---

*"I think it can make life easier for humans as tasks can be performed faster. If the AI is made/programmed well then it will also eliminate human error meaning that the tasks are performed more efficiently" - Workshop participant, London*

---

During the workshops we shared a demonstration of ChatGPT. The belief that AI would create outputs in a faster and more efficient way was often reinforced during the live demonstration of Chat-GPT during the workshops. Participants who had not used Chat-GPT or other generative AI tools before tended to be surprised at the speed at which Chat-GPT created work and the quality of the work it produced.

## 2.4 Spontaneous concerns about the application of AI technology



Figure 3: Word cloud summarising responses to the question asked in an online survey shared prior to the workshop: "What do you see as key drawbacks of artificial intelligence?"

Amongst general public participants, concerns about the use of AI technology tended to focus primarily on the potential for increasing use of AI to lead to job losses. As

<sup>3</sup> <https://www.turing.ac.uk/news/publications/how-do-people-feel-about-ai>

participants assumed that AI will deliver work faster and more efficiently than humans, there is an assumption that many people may find that they are replaced in their jobs by AI. Again, this reflects finding from recent Thinks Insight & Strategy polling where 67% said that they thought increasing AI use would lead to more jobs being lost than created, with those aged 45-64 (73%) and disabled people (74%) particularly likely to be worried about job loss.<sup>4</sup> Only 13% believe that AI would lead to more jobs being created than lost.

Additionally, many participants raised concerns about the potential for humans to become overly dependent on AI. There was a sense that this could lead to people neglecting to develop key skills as they were used to AI completing those tasks on their behalf. These participants often referenced examples from their own lives about skills they have lost as a result of technological improvement (for example, struggling to navigate using a paper map as a result of typically using Google Maps to get to places). This reflects findings from The Turing Institute which identified that the public most commonly express concerns about human overreliance on AI technology over professional human judgement<sup>5</sup>.

Finally, a smaller (but still sizeable) number of participants raised concerns that increasing use of AI would lead to a loss of empathy in decision-making and communications. This was a result of broader belief that AI would make decisions in a 'cold', rational way and would struggle to communicate these in an empathetic way.

## 2.5 Spontaneous views of UK Government use of AI

Participants across all strands spontaneously believed that AI will be an important technology in both the short- and long-term, with high levels of uptake across all sectors. As such, there was a widespread and spontaneous belief that AI will 'inevitably' be used by government departments for a wide range of tasks.

---

*"If it's implemented in a credible and robust way - there can be many uses [for AI].  
AI has been a great time-saver for us and can do the same for DfT" - Stakeholder*

*"I prefer human interaction however if the AI is developed well, I could really see  
this taking off" - Workshop participant, London*

---

Despite the belief that AI will be an important (and likely widespread) technology in coming years, participants tended to not express strong positive or negative views about the potential use of AI by the UK Government. In large part this appeared to be a result of an assumption that AI usage was not personally relevant to them (for example, participants were worried about job losses within society, but few were personally worried about losing their jobs).

Alongside this, there was a broader sense that opposing AI use was likely to be futile, with the speed and efficiency benefits of AI use likely to be impossible for organisations (such as the UK Government) to ignore. Some participants across all locations emphasised the

---

<sup>4</sup> Ibid.

<sup>5</sup> <https://www.turing.ac.uk/news/publications/how-do-people-feel-about-ai>

importance of the UK Government keeping up with the pace of change regarding new technologies such as AI.

## **2.6 Spontaneous views of potential uses of AI within UK Government**

Participants tended to believe that the best uses of AI within UK Government would involve using it for either analytical or administrative tasks. Participants felt these kinds of tasks both played to AI's strengths (i.e. speed and efficiency, computing and calculating) and were areas where humans were more likely to make mistakes (for example, as a result of poor attention to detail). Examples of the kinds of tasks participants felt UK Government could use AI for were: speeding up or automating repetitive and mundane processes such as: performing repetitive tasks or basic administrative duties, monitoring and organising databases, verifying documentation, conducting advanced statistical calculations, and accounting and budgetary tasks.

Importantly, none of our participants spontaneously raised the idea of using AI for correspondence or communication related tasks. This was likely due to the belief (covered in 2.4) that AI is unlikely to be good at communicating in an empathetic way with humans.

## 3. Overarching views of the use cases

### 3.1 Comfort with the use cases

#### 3.1.1. Overarching comfort

Across both the workshops and the 'engaged' citizen interviews, general public participants tended to say they were comfortable with both use cases (especially once DfT's proposed mitigations had been presented to them), although levels of comfort varied between them. Participants were consistently more comfortable with the use of AI for the analysis of consultation responses than with the use of AI for correspondence.

This difference in levels of comfort reflected the spontaneous associations and assumptions that participants had about what AI would be good and bad at (covered in more detail in section 2). Participants assumed that AI will be better at large scale data analysis than humans, meaning that the idea that AI could be used for analysing consultation responses felt intuitive to participants and was generally accepted. Conversely, participants were sceptical about the ability of AI to act in an empathetic manner, meaning many felt it would be unsuited to writing correspondence. These findings are in line with those from The Turing Institute which showed the public are more likely to be concerned with AI technology that seeks to replace human judgement and decision making<sup>6</sup>.

Finally, concerns around potential job loss continued to be an important factor in overall levels of comfort about the two use cases. Participants across the sessions believed that using AI for these use cases would almost certainly mean job losses in the future.

#### 3.1.2. Sub-group differences

In addition to differences in comfort between the use cases, there were also differences in overall comfort in different locations and within some sub-groups within the public.

There were participants across the locations who were more sceptical about the use cases. However, participants in London and Newport tended to be more comfortable with both use cases than participants in Glasgow. Whilst participants in all locations raised

---

<sup>6</sup> <https://www.turing.ac.uk/news/publications/how-do-people-feel-about-ai>

similar views of the benefits and drawbacks of AI, participants in Glasgow tended to feel less strongly about the benefits and feel greater concern about the drawbacks. This was particularly the case when it came to concerns around job loss (for both use cases) and concerns around the ability of AI to write empathetically (for correspondence).<sup>7</sup>

---

*"I am still uncomfortable with the way this would improve efficiency as there is still cost and need for human interaction to get success" - Workshop participant, Glasgow*

---

Alongside this, participants across workshops and interviews who identified as being more sceptical or fearful of technology within the sessions (i.e. self-describing as 'technophobic' or similar) were generally less comfortable about both use cases.

---

*"I am a bit concerned about this technology and it would take a lot more to develop if its going to do a good job. The sheer amount of variability in the data provided will make it difficult to process" - Workshop participant, Newport*

---

These were often, but not always, older participants who expressed concern that the increasing use of technology will create access issues for those who are less confident using technology. These participants were also often less confident about whether AI will produce accurate or high quality outputs at its current stage of development.

## 3.2 Factors affecting views of the use cases

When discussing both use cases with general public participants we saw four key factors influence responses. These were the salience of the use case to participants, the personal relevance of the potential benefits and drawbacks, the perceptions of the quality and accuracy of AI for the specific use case and views of human oversight. In the section below we cover each of these areas in detail, including an overview of participants' views of each and how they influenced their responses.

### 3.2.1 Salience of the use case

Views of the use cases were heavily affected by how salient or otherwise participants found them.

Most workshop participants had little to no direct experience of either use case and found both to be relatively low salience as a consequence. However, participants tended to find use case 2 (correspondence) more personally salient than use case 1 (consultations). This was because, whilst they had not had experience of sending correspondence to a UK Government department, most had some experience of corresponding with organisations/businesses and dealing with customer service. This experience was seen by most as comparable and thus informed their views of the use case.

---

<sup>7</sup> Please note that as this was a qualitative study we can't definitively say if this finding is generalisable to a broader population (e.g. whether those in Scotland are more sceptical than those in England or Wales).

Conversely, the 'engaged citizen' participants all had some experience of either responding to consultations or correspondence with government (either UK or at a local level). These participants tended to find both use cases more salient than the workshop participants and their experience of both informed their responses.

### 3.2.2. Personal relevance of the potential benefits and drawbacks

When discussing the use cases general public participants tended to find benefits and risks that impact on the public more compelling and resonant than those which would affect DfT. This sentiment is echoed in research for the Centre for Data Ethics and Innovation which found that the public are more comfortable with applications of AI in the public sector which benefit the public, rather than 'just' public sector workers<sup>8</sup>.

In line with this, the most resonant benefits for both use cases related to the increased speed of response (both to consultations and correspondence) that AI could deliver. The most resonant drawbacks related to the quality of outputs the public would receive (i.e. will they be accurate or empathetic).

---

*“Responses can be dealt with quickly and efficiently.” - Workshop participant, London*

*“AI lacks emotional intelligence to handle complex situations.” - Workshop participant, Glasgow*

---

Conversely, the least resonant benefits were those that related to 'efficiency' (both time and money) for DfT, which was seen to have limited impact on participants as members of the public. Importantly, there was suspicion from many participants that the term 'efficiency' was being used euphemistically, and would ultimately mean job losses for people working in the Department. Similarly, the cost of creating the AI was not a particularly resonant drawback for participants.

Given the tendency of participants to focus on how use cases will directly impact the public rather than how they will impact the Department it will be important for communications on the use of AI to focus on what it means for the public. It cannot be taken for granted that the public will intuitively understand (or trust) that more indirect benefits as a result of time and cost efficiencies will ultimately benefit them.

### 3.2.3. Perceptions of quality and accuracy

Whilst few participants raised quality or accuracy of outputs as a potential concern when discussing AI use overall, accuracy and quality became major focuses of discussion when talking about specific use cases.

Most workshop participants had little (and more often no) personal experience of responding to consultations or correspondence with UK Government departments. For many the most front-of-mind examples of interacting with AI technology were interacting

---

<sup>8</sup> <https://www.gov.uk/government/publications/public-perceptions-towards-the-use-of-foundation-models-in-the-public-sector>

with AI customer service chat bots, which many described as being a negative and frustrating experience. These negative experiences often activated concerns around quality and accuracy, with participants raising examples of where their query or problem had been misunderstood by AI, leading to assumptions that the AI DfT develops may also have similar problems.

---

*"Bring back the human touch and empathy otherwise the complaints will increase due to frustration of not being answered or acknowledged" - Workshop participant, Newport*

---

This focus suggests that public support for the use of AI by DfT is contingent on it developing AI that delivers high quality outputs. If it does not deliver outputs of this kind of quality then the public are unlikely to believe that the broader benefits of speed and efficiency will be truly realised. This is also reflected in research conducted for the CDEI which found that the benefit of increased speed is of no use if AI produces inaccurate outputs<sup>9</sup>.

### 3.2.4. Views of human oversight

General public participants felt it was essential that humans have oversight of AI outputs in both the short-term and longer-term. For many their acceptance and comfort with the use of AI in these use cases was contingent on retaining human involvement. This is also reflected in research *Thinks Insight & Strategy* conducted for the Centre for Data Ethics and Innovation which showed that public comfort with the use of AI in the public sector was contingent on human oversight of the process<sup>10</sup>. This was due to two key reasons, quality assurance and accountability.

Regarding quality assurance, participants felt that human oversight would be crucial to ensure that AI outputs were of a good standard, as well as for training and monitoring the overall performance of the AI. This was seen to be particularly important for when AI will be asked to answer complex or emotive queries which participants assumed AI would be more likely to struggle with.

---

*"It will be reviewed by a human and hopefully they can add personalisation and emotion if needed." - Workshop participant, Glasgow*

---

In relation to accountability, participants felt that human oversight would be crucial for maintaining trust in the technology. There was a strong belief that AI would be unable to be accountable for its outputs and that it is important that someone takes responsibility for its outputs, especially in the case where there are inaccuracies.

---

<sup>9</sup> <https://www.gov.uk/government/publications/public-perceptions-towards-the-use-of-foundation-models-in-the-public-sector>

<sup>10</sup> <https://www.gov.uk/government/publications/public-perceptions-towards-the-use-of-foundation-models-in-the-public-sector>

---

*"I prefer if there was some human interaction at some level because AI is prone to misinformation." - Workshop participant, London*

*"There's a potential for error due to lack of common sense - humans needed to check the work of AI". - Workshop participant, London*

---

However, whilst human oversight was seen as crucial, there were some concerns about the risk of human complacency when it comes to reviewing AI outputs. Consistent with the concern that humans may become overdependent on AI, there is a belief that civil servants may become overconfident in the quality of AI outputs and fail to thoroughly review them, leading to potential inaccuracies in outputs not being caught. Alongside this, a smaller number of participants had concerns that an overreliance on technology may lead to the 'deskilling' of workers, who will not have sufficient experience of creating outputs themselves to effectively review them.

---

*"There should be restrictions in place to prevent overdependence on AI. Human and AI should be used interchangeably" - Workshop participant, London*

---

A minority of participants, however, flagged that vulnerabilities around poor quality checking was likely to already exist within the current system. These participants were consequently less worried about the potential for civil servants to fail to sufficiently quality check AI generated outputs.

### 3.3 Stakeholder views

Overall stakeholders were comfortable with both use cases. Stakeholders often had some experience of, or had been privy to, discussions within their own organisations about using AI. This meant that they often had higher levels of knowledge about AI and its capabilities and therefore often had fewer concerns about quality and accuracy than the general public. However, despite their overall comfort with both use cases, stakeholders tended to be sceptical about the extent of the benefits of the use cases, with some questioning the extent to which AI will address the key issues affecting both consultations and correspondence.

Some stakeholders felt that the key cause of the issues DfT faces with analysing consultations relate more to the design and execution of consultations rather than the volume of responses and their analysis. These stakeholders believed that consultations 'over rely' on open questions and are insufficiently tailored to respondents, meaning that consultations gather large amounts of complex and ambiguous responses.

Alongside this, other stakeholders queried whether delays in publication of responses to consultations was the result of the need to undertake large amounts of analysis. These stakeholders tended to believe the key issue that led to delays in publication of responses was what they termed the need for 'political interpretation' of consultation responses, especially when consultations results challenge the course of action the Department would like to take. These stakeholders were therefore sceptical about the extent to which speeding up analysis (through the use of AI) would speed up consultation response times.

---

*“We don’t have a problem with the way they respond to the consultations. Often the delay is about Ministerial decisions and what the findings of the consultations are about.” - Stakeholder*

---

In relation to correspondence, some stakeholders felt that the (unspoken) goal of responses is to shut down, rather than open up, conversation. They believe this is due to broader concerns about freedom of information requests. These stakeholders felt that responses to correspondence are already effective at delivering these limited responses. As a result, they felt that AI would likely be deployed in a use case where there was limited need and, at worst, could be interpreted by stakeholders as a clear signal that there is limited interest in their correspondence.

---

*“I think there is a challenge to the democratic mandate if you feel you’re getting machine responses rather than human responses” - Stakeholder*

---

## 4. Perspectives on Use Case 1: Using AI to analyse consultation responses

### 4.1 Public response to use case 1

#### 4.1.1. Levels of comfort

Overall, both the interviews and workshops showed that the public generally felt comfortable with the use of AI to analyse consultation responses for internal use. The only exceptions were in keeping with the sub-group differences described in 3.1.2. That being said, these differences were rather minor in the case of consultation analysis, and the majority of participants across the board felt relatively comfortable with the application of AI in this case.

#### 4.1.2. Drivers of comfort for the use of AI to analyse consultation responses

Comfort with the use of AI to analyse consultation responses was driven by three factors: belief that data analysis is a task that AI will be good at; low expectations of a personalised response to consultations; and satisfaction with the proposed level of human oversight.

As mentioned in section 2.6. there was a spontaneous assumption from many participants that AI will be able to perform data analysis faster and more effectively than a human.

---

*"[AI] can make information easier to digest for whoever interprets the data" - Workshop participant, Newport*

*"AI could reduce human error by sorting the data, when they're grouping responses together and removing duplicate responses...AI could be more reliable." - Workshop participant, London*

---

Alongside this, the nature of responding to a consultation meant that many participants felt comfortable with the idea of AI being used for data analysis. Given consultations involve large numbers of responses, there is an expectation responses from the public will be aggregated in order to form a larger picture of the public's opinions on a given issue. As

such, there is no expectation that respondents will receive a personalised response. Consequently, few participants were concerned about the potential loss of a human touch that many fear might come with increasing use of AI (covered in more detail in section 2.4.).

Finally, participants were generally satisfied with the level of human oversight for this use case. The fact that AI would perform the data analysis but humans would still be responsible for crafting the final consultation response, meant that the process still allowed humans to perform a critically authoritative part of the process. This also meant that if something went awry, members of staff would at best be able to notice and rectify it, or at the very least be someone who could be held accountable for the outputs and outcomes.

---

*"Just quick accurate and impartial results. It will be ok'd by a human at the end as well so it's all good" - Workshop participant, Glasgow*

---

### 4.1.3. Drivers of discomfort for the use of AI to analyse consultation responses

Despite there being high levels of support overall, general public participants did highlight some concerns about the use of AI for analysing consultation responses.

#### 4.1.3.1 Accuracy

The most commonly raised concern by participants related to accuracy. Some participants were concerned that technology may not produce fully accurate outputs, drawing on their experience of other AI powered services misinterpreting open text responses. For example, many had used AI customer service chatbots in the past which had misinterpreted their requests and given inaccurate responses. This led them to assume that AI may not always be able to understand open text responses and would therefore struggle to accurately analyse them.

During discussions, participants raised the following as types of responses they felt AI may struggle to interpret: longer and more complex responses (i.e. when a single response makes multiple points or touches on a variety of issues); responses making use of rhetorical techniques such as sarcasm or hyperbole; responses using slang and responses with spelling mistakes or written in poor English.

While many shared these concerns, other participants noted that humans were also likely to struggle in the face of these challenges. They, too, may struggle to consistently analyse these kinds of responses thoroughly, and are also likely to make simple mistakes in attending to details, especially given the vast volume of responses which some consultations receive.

---

*"If people respond with incorrect grammar it may not understand what they are trying to say" - Workshop participant, London*

*"A drawback is AI not doing its job very well which will require more reading and correcting. It's going to take time and money, it's an investment for a really good pay off"*  
*- Workshop participant, Newport*

*"Sure, you can have key words here and there, but humans aren't good at using the exact same words every time" - Workshop participant, Newport*

---

#### 4.1.3.2 Bias

Another issue that was raised was the potential for bias. Once participants understood the way in which AI models are developed, some were concerned about the potential for poor quality input data being built into the system leading to biased results - risking marginalised groups being disadvantaged. Participants' knowledge of how AI might produce biased outputs was limited, so discussions of what should be done to mitigate this risk tended to focus on a more generic desire for DfT to ensure the use of high quality input data for the training and development of the AI model.

---

*"It would have a human preference, presumably meaning bias" - Workshop participant, Glasgow*

*"If data is input incorrectly at the beginning will be a problem" - Workshop participant, London*

*"I would prefer if there was some human interaction at some level because AI is prone to misinformation" - Workshop participant, London"*

---

#### 4.1.3.3. Job loss

Another notable concern was potential job loss. A few participants were quite concerned about whether the implementation of AI technology would mean significant job losses or redundancy within DfT both in the short- and long-term. Many were keen to ensure that staff would not be "replaced by robots", and that if needed they would be redeployed to other roles within the Department. These concerns were met by more optimistic views that whilst some jobs may be lost, new ones would be created over time.

---

*"Cutting out the middle man could result in job loss" - Workshop participant, Glasgow*

---

#### 4.1.4. Responses to DfT's proposed mitigations

During the workshops and interview, the participants were presented with several mitigations to tackle the potential challenges that could arise. These included a human staying 'in the loop', optimising the model to minimise issues with accuracy and curating the training data to allow AI to understand different types of language. Discussions showed that most participants believe the mitigations were sufficient. Moreover, the fact that human oversight was still a key component of the process reassured them that any mistakes made by the AI would be identified and fixed. Additionally, the fact that humans would still be the ones to write up and present the data helped preserve the important 'human touch' that many participants spoke about. This served to assuage concerns about

potential errors in the process, and contributed to participants' overall levels of relatively high comfort with this use case.

---

*"Humans will still be in charge of the outcomes, so it feels safe" - Workshop participant, Newport*

---

That being said, some felt that even more could be done to mitigate potential issues.

---

*"Are they still going to be gathering information from humans, local people?" - Workshop participant, Glasgow*

---

These ideas included:

- (a) Conducting comparative analyses to assess quality. This would entail piloting the AI and comparing its outputs to human-led analysis on the same data
- (b) Quality controls on model development and training data.
- (c) Establishing a regulatory system for AI, either both on a national level or within the Department. This was of particular interest to participants who were concerned about data security.
- (d) Ensuring consistent and ongoing updating of AI training data, so that it can keep up to date with developments and changes in language.

## 4.2. Stakeholder response to Use Case 1

Stakeholders were broadly comfortable about the use of AI to analyse consultation responses but had more concerns than the general public. They felt comfortable with AI across the board and appreciated the advantage of faster analysis associated with AI believing it has the potential to make the consultation process quicker and more efficient. That said, they also had concerns about bias and the extent to which AI would be able to accurately understand views expressed from organisations and individuals. For example, the technology may disregard an important response if it was only made once or fail to understand the policy context behind the response.

---

*"This seems like a fairly typical use and with the volume [of data] it makes sense." - Stakeholder*

*"Context is king! We choose our words carefully in responding and usually it takes an experienced Government official to know exactly what we're saying – could a machine be trained to read between the lines?" - Stakeholder*

---

In addition, some stakeholders expressed reservations as to how relevant and impactful this increased speed and efficiency will be to them and the issues they write in about when responding to consultations.

Stakeholders understand consultation analysis to be a time-consuming process, and thus do not generally expect (or require) results to be published quickly. When stakeholders prepare a response, they rather prioritise quality of analysis over speed. The expectation is instead that their response ultimately be ‘properly’ considered by a human.

In this context, stakeholders emphasise the technology must be rigorously piloted to ensure it is capable of delivering a similar quality response to humans.

---

*“I spend a lot of time and effort in getting information which I want to help shape Government policy... the hope is that somebody will think about it.” - Stakeholder*

*“If you know what the rules are, then I think you're more understanding about how some things take a little longer.” - Stakeholder*

---

Moreover, they hold some scepticism about whether the analysis is what slows down the publishing of results, rather than the internal politics of it (i.e. when consultation responses don't give the answer that UK Government ‘wants’ to hear).

---

*“We don't have a problem with the way they respond to consultations [at the moment]. Some take too long to come out, but from what we have gathered it's not because they're too complex. Often the delay is about ministerial decisions and the findings of the consultation.” - Stakeholder*

---

## 5. Perspectives on Use Case 2: Using AI to respond to correspondence

### 5.1 Public response to Use Case 2

#### 5.1.1. Levels of comfort

The use of AI for correspondence is more divisive with general public participants than its use for consultations. Whilst some participants were very comfortable with the idea of using AI to analyse and draft external correspondence, others were strongly opposed, and it was in this use case that the sub-group differences described in 3.1.2 came through most strongly. Participants across the groups who were hesitant of new technologies generally expressed stronger doubts and were less persuaded by the potential benefits of AI when applied to use case 2. On the other hand, those who were more comfortable with technology in general felt more comfortable with this use case.

---

*"You shouldn't use AI for correspondence because AI lacks emotional intelligence to handle complex situation. AI won't send your issues to specific people, instead it sends to teams" - Workshop participant, Glasgow*

*"It's cost effective. You can receive your response much quicker than if it was done all by humans, it is unbiased. It is accurate and the costs saved can be used elsewhere in the business" - Workshop participant, London*

---

#### 5.1.2. Drivers of comfort for the use of AI to respond to correspondence

When it came to using AI to respond to correspondence, participants expressed mixed reactions. On the one hand, there were a few factors that participants noted as having the potential to be beneficial and reassuring. This included the clear view that AI would be able to deliver speed, producing results much faster than the current 20-day timeline. Participants highlighted this as advantageous to the general public, pointing to the fact that they would receive quicker responses to their letter or email. Additionally, similar to use case 1, the public were greatly comforted here as well by the fact that AI-generated responses would be reviewed by a human before sending.

---

*"You will get a quicker response from DfT if they use AI and they assure that every response is checked by a human" - Workshop participant, London*

---

### 5.1.3. Drivers of discomfort for the use of AI to respond to correspondence

On the other hand, the public had greater concern about using AI to respond to correspondence than use case 1. Participants believed that in order for the public's experience of correspondence with DfT to be positive, it need to not only be timely but also accurate and empathetic, and there were questions as to whether AI would be able to deliver this.

---

*"I think if someone has taken the time to correspond then they do not want a generic response" - Workshop participant, Newport*

---

#### 5.1.3.1. Emotional intelligence

The public reiterated a worry that AI's lack of emotional intelligence would result in its crafting generic and unempathetic replies to what may often be emotional, personal enquiries regarding sensitive issues—leaving respondents feeling that they have not been considered or respected. This sentiment is also reflected in research for the CDEI which showed the public do not believe AI has the emotional intelligence to make 'good decisions' about individuals<sup>11</sup>.

---

*"AI might not understand your explanation. I wouldn't feel valued, but [it's] rather just another impersonal letter" - Workshop participant, Newport*

---

#### 5.1.3.2. Accuracy

There were once again concerns as to whether AI would give factually accurate responses reliably. Participants expressed doubt that AI would be able to properly make sense of and resolve many queries, especially more complex or niche ones. They understood correspondence to touch on a significant breadth and diversity of topics, communicated in any number of ways, and were sceptical of AI's ability to handle them all appropriately.

---

*"Feel the lack of emotion and knowledge will mean its not a good decision and give poor response because its based on previous correspondence" - Workshop participant, Glasgow*

---

---

<sup>11</sup> <https://www.gov.uk/government/publications/public-perceptions-towards-the-use-of-foundation-models-in-the-public-sector>

### 5.1.3.3. Job loss

And finally, consistent again with use case 1, some participants described a fear that adopting this technology would result in job loss within the Department.

It is noteworthy though, as evidenced in the above, that while use case 2 elicited more mixed reactions and generally higher levels of concern than use case 1, it did not present a different set of concerns—the worries were the same, they were just experienced and expressed more frequently and strongly among the general public with this scenario.

---

*"Make sure the tech doesn't make people redundant, don't want to lose jobs" - Engaged public citizen*

---

### 5.1.4. Influence of previous experiences on views

Perceptions of the potential benefits and drawbacks of this use case were heavily influenced by previous experience interacting with AI in customer service contexts. Participants recalled customer service interactions they had experienced across other organisations and services, the nature of which then shaped their assumptions of what an interaction with AI would look like when applied to this context.

Participants repeatedly emphasised the importance of feeling like they were 'being heard' by the organisation they were reaching out to in these situations. Many felt this was of the utmost importance in creating a positive customer service interaction, and many described frustrating past experiences engaging with chatbots and other AI-assisted technology that failed to live up to this standard. These negative experiences included:

- Receiving inappropriate/inaccurate responses
- A failure on the part of AI to address their core issue/problem
- Not being given the option to speak to a human

These negative experiences coloured their views of the potential merits of AI in this use case, increasing scepticism. They expected that if AI did not address their problems in these other past customer service interactions, there was little likelihood that this time and this application would be different.

---

*"[AI] doesn't understand [you] sometimes and it really drives you crazy." - Workshop participant, London*

---

However, a minority of participants who regularly use generative AI to draft their own correspondence were more optimistic. They referenced the skill of ChatGPT and suggested that an even better trained tool should generate satisfactory correspondence. Drawing on their own more satisfactory interactions with AI-assisted technology in the past led them to believe that when applied to this use case, it could deliver similar benefits and positive outcomes.

Finally, a small number of participants with experience of working in customer service roles said that, in their experience, emails were always handled with 'stock' responses and rarely tailored to individuals. These participants assumed that this is what currently happens with correspondence to AI now and therefore believed that, as long as AI is responding quickly, this would represent an improvement.

#### 5.1.4. Responses to the mitigations

DfT's proposed mitigations mostly addressed concerns about AI's ability to write empathetically and provide a satisfactory customer service experience.

---

*"Humans will check every response though generated by AI to reassure you" - Workshop participant, London*

---

When presented with various mitigation strategies, participants largely focused on the approach of maintaining human involvement. They were greatly comforted by the fact that a human would continue to fact-check all correspondence for the following reasons:

- (a) Quality control: The fact-checking process would act as a fail-safe around questions of AI responding accurately, particularly for those complex or niche concerns.
- (b) Empathy: A reviewer also prevented the all-important "human touch" from being lost entirely (which is noted as particularly crucial for sensitive matters).
- (c) Accountability: Continuing human involvement would mean there would always be someone who could be accountable for any issues should they arise.
- (d) Job preservation: Maintaining staff engagement in the process helped allay some fears about potential redundancies.
- (e) Trust: Preserving a central role for humans could help the public establish and maintain trust in the new system/technology.

#### 5.1.5. Responses to labelling approaches

To evaluate various communication styles, after discussing this use case participants were presented with three versions of prototypical language indicating that a response was written using AI. Given the focus and premium participants placed on continued human involvement, it is unsurprising that they preferred labelling that explicitly articulates it—finding those messages to be the most transparent and reassuring. A summary of their overarching perceptions of each message is below:

<p>1. 'To improve efficiency, this response was drafted with assistance from AI. It has been fully checked by a human.'</p>	<p>This message received the most overwhelmingly positive response.</p> <p>This was largely because it both offered an accurate explanation of the process while</p>
---	--

	clearly highlighting the role of the human—a key focus for participants.
2. 'Developed using AI and reviewed by humans.'	This message was received less positively than the first.  While still referencing human involvement, this message was seen to be rather cold and did not offer enough contextual information to be deemed valuable.
3. 'Crafted using AI technology.'	Overall, this message was not well received.  It was the least transparent and importantly did not communicate that any human review had taken place. There was seen to be some potential for this message to aggravate feelings of frustration for those sending more emotional or personal messages and creating a sense that those messages have not been read at any stage by a human.

It is important to note that whilst most participants did believe that DfT should label correspondence drafted by AI in their response, this was not the view of all participants. Some questioned whether it was relevant or important and others said that, given how sceptical they felt others were about using AI to respond to correspondence, explicitly mentioning that AI has been used to draft a response might risk undermining trust in the response, even if it was correct.

## 5.2 Stakeholder response to Use Case 2

Similar to the general public, stakeholders again had mixed views of the use of AI for responding to correspondence, with some expressing comfort with the idea and others opposing it.

Stakeholders saw value in AI use for increasing speed and efficiency of processing 'standard' correspondence. They also imagined that adopting this technology would take a heavy burden of mundane and repetitive tasks off staff. They understood the benefit of improving consistency and potentially even quality of responses, should the AI be trained to a high level, and like the public, they were reassured that a human reviewer would remain involved.

---

*"This will benefit the Department as it will take away the drudgery for the staff...[and] the public after they start to build trust in the system." - Stakeholder*

---

Stakeholders, too, expressed drawbacks of AI adoption for use case 2. They hypothesised that if AI sends standardised and unempathetic responses, it may seem as though DfT

'doesn't care' about the public they are meant to serve. This would paint the Department in a negative light and damage their reputation—a concern stakeholders highlighted as being an important one to factor into and address during development and roll-out.

Additionally, while they appreciated the role of a human reviewer, there remained questions as to whether those human reviewers would be thorough enough in their checks to catch any biases in data inputs or outputs—presenting the possibility that these biases would be integrated into the technology and further perpetuated.

---

*“There are going to be opportunities to make things easier but there needs to be a boundary between AI helping the process but not driving it!” - Stakeholder*

---

And finally, stakeholders noted that outputs could be high quality if the tool were to be trained properly, but equally acknowledged there were considerable risks if such standards weren't met. Namely, they expressed concern as to whether the technology would be able to properly answer questions consistently and stay up-to-date on policy changes.

## 6. Guiding principles of the application of AI

### 6.1 Guiding principles

As outlined in the research objectives, one of the key aims of this project was to give participants the opportunity to develop guiding principles for the application of AI within DfT. These principles would encapsulate the issues that were most important to the research participants and share the ways in which they imagined they could be addressed such that it would benefit the wider public. Ultimately, the research yielded six core guiding principles that DfT could consult as they move forward with the development and rolling out of AI technology within the context of the two use cases. The principles are as follows:

#### 6.1.1. Ensure the technology works (i.e. is highly accurate) before using it.

General public participants were adamant that AI technology would only be beneficial if it was accurate and did the job as well as, if not better than, humans. They believed it was paramount that a switch to AI technology does not result in a drop in quality, and ideally, improves it. Many reiterated that time saving and other associated benefits would be irrelevant if the quality of outputs were not up to standard. Some stakeholders also shared similar views, noting that inaccuracies could lead to reputational damage for DfT and loss of trust in the Department.

---

*"Thoroughly test and check AI - have proof this it works properly, proper maintenance, regular updates and development" - Workshop participant, Newport*

---

#### 6.1.2. Start small and roll out the technology slowly

In order to deliver the first principle many participants, particularly those who were a bit wary of the use of AI, stated that it was important that the technology is introduced slowly in order to ascertain its accuracy and effectiveness. This includes conducting several pilot tests of the AI, including those geared towards comparing AI outputs to human generated outputs to ensure comparable quality. A slow roll out would also give the Department the opportunity to fix any glitches before rolling it out more widely, as well as give the public time to adjust to the change.

---

*"Continuous evolution - not set it and forget, slower roll-out, start small, gradually increase" - Workshop participant, Newport*

---

### **6.1.3. Keep quality front and centre**

Several participants and stakeholders echoed the importance of consistently monitoring the AI and gathering continuous feedback about what is and is not working. As such, the service would not only remain up-to-date with the latest policy changes, but continue to evolve and further improve as time goes on.

---

*"Quality of AI should be maintained and regularly improved" - Workshop participant, London*

---

### **6.1.4. Retain human involvement to ensure quality and that there is accountable if something goes wrong**

Many participants and stakeholders felt that humans represented an important level of quality assurance in the process. For most, the value of human oversight was quite high and having a human involved allowed them to trust both the technology and the overall process. It was also very important that a person, rather than a machine, would be responsible for outputs and answerable if need be when issues arise.

---

*"AI should flag when it doesn't understand something so a human can review"*  
*- Workshop participant, London*

---

### **6.1.5. Consider excluding complex cases from AI**

Many participants suggested that parameters be put in place for the types of issues that could be addressed using AI technology. Many suggested that certain cases which may be more niche, complex, or require more empathy may be better suited for humans to respond to, rather than AI.

---

*"Filtering by complexity/simplicity of issues" - Workshop participant, Newport*

---

### **6.1.6. Be transparent about the use of AI**

There was a lot of discussion around the need for transparency in the application of AI. Participants suggested that DfT clearly communicate the ways in which they intend to use the technology, the rationale for doing so, and the benefits it provides, so the public could understand the value that the technology offers both to the Department as well as the general public.

---

*"Tell people why - give clear examples. i.e, 'we've got so many responses and so many people'" - Workshop participant, Newport*

---