# Summary Report

## Public dialogue on AI in local government: interim results

ai@cam | UNIVERSITY OF CAMBRIDGE

Hopkins Van Mil

# Contents

# Summary

What does the public want from AI in local government, and what does this mean for those developing it? This interim report synthesises findings from the first phase of a public dialogue with Cambridgeshire residents about AI in local government. The dialogue suggests cautious optimism: residents can see significant potential for AI to improve services, but their support comes with clear conditions. For researchers, developers and local authorities working in this space, the findings point to several priorities:

## Summary

### ❖ Start with the 'right' problems

The public wants AI R&D to focus on areas where real needs exist and where public benefits will be greatest, demonstrating the ability of AI technologies to be used in local government to improve quality of life. Innovation for its own sake, or deployment driven primarily by cost-saving opportunities, will not gain public trust or support.

### ❖ Start with the 'right' problems

R&D efforts should focus on applications where AI does what machines do well - routine tasks, data synthesis, administrative process - freeing humans to apply empathy, judgment and creativity. The public does not want AI making decisions that directly affect people's lives, especially vulnerable populations. This means designing systems with meaningful human oversight built in from the start, not added as an afterthought.

### ❖ Prioritise proven effectiveness.

Residents are frustrated with poorly functioning digital systems that seem designed to cut costs rather than improve service. They want to see AI systems thoroughly tested and proven to provide reliable, high-quality service, learning from what works in practice before scaling up.

### ❖ Design for inclusion from the outset.

R&D must account for digital exclusion, accessibility needs, and the diversity of residents' capabilities and circumstances. This means always maintaining human interaction options, designing for those with lower digital literacy, and ensuring systems work for everyone.

### ❖ Build trust through transparency and accountability.

The public wants to know how systems work, who developed them, what data is being used, and where savings are being reinvested. R&D efforts should include clear mechanisms for explainability, monitoring, and independent evaluation from the start, including named people repsonsible for the use of AI.

### ❖ Plan for an uncertain future.

Implementation plans should include robust contingency planning, for example: exit strategies if systems fail, maintaining staff capabilities to work without AI, and security measures against both malicious attacks and system failures. This risk-aware approach should be fundamental to design, not an afterthought.

### ❖ Involve the public throughout.

Residents want to be engaged at every stage - from setting principles and priorities, through design and testing, to ongoing monitoring and improvement. R&D efforts that fail to meaningfully involve diverse service users and members of the public risk developing solutions that don't meet real needs or gain public acceptance.

### ❖ Consider the broader context.

AI deployment doesn't happen in isolation. The public expects consideration of climate impacts, job transitions and training needs, data security implications, and how AI fits with councils' broader goals like Net Zero commitments. R&D efforts need to account for these systemic factors, recognising both these wider interactions and the social and economic context for deployment.

**Ultimately, these interim results show that residents are willing partners in developing AI for local government - but only if it serves their needs, protects what they value about public services, and is developed with their involvement.**

# Introduction

In October 2025 ai@cam commissioned the social research agency Hopkins Van Mil (HVM) to design and deliver a public dialogue on AI in local government.

# Introduction

ai@cam is the University of Cambridge's flagship mission on AI, which is driving a new wave of AI innovation that serves science, citizens, and society. ai@cam provides an infrastructure for interdisciplinary AI innovation through research, education, and innovation programmes that leverage the world-leading research taking place across the University of Cambridge and connect this to real-world needs.

Hopkins Van Mil is a social research agency specialising in deliberative processes which bring people together to explore and understand society's challenges. For over twenty years we have designed and facilitated public dialogues including people across society in open and constructive conversations to build mutual respect and understanding. Our work which includes people across society, leads to actional insights, collaborative solutions and evidence-based policy making.

**The aims of this deliberative process were to engage and work with a robust sample of ~100 participants from Cambridgeshire to:**

- Help ensure research and development activities in AI in local government services connects with public values, expectations and needs;
- Develop understanding of 'why' people's principles and conditions for AI in local government are important (or what these conditions look like in practice for R&D activities);
- Inform the selection criteria for future research projects in AI for local government convened by ai@cam;
- Explore in-depth specific AI in local government use cases e.g. for adult social care;  council planning; customer contact centres; transport and transport infrastructure and homelessness prevention.

**Through the dialogue process running from November 2025 to March 2026 key questions are being explored such as:**

- What are participants' hopes and concerns for AI in local government?
- What are the conditions, caveats and red lines they apply to AI in local government?
- What would make AI in local government best align with public needs, aspirations, values and principles?
- How can we understand more about why people feel as they do about AI in local government?
- How can research institutions align AI projects with public views and values?

ai@cam


# Method

**The first phase of this public dialogue is now complete. A recruitment specification was devised to recruit ~100 people from Cambridgeshire, broadly reflective of the UK population. 95 people joined the process and were supported to take part in activities and discussions in a tailored online space[1] for the initial part of the dialogue. For five days participants:**

- Reviewed contextual videos on what public dialogue is; the purpose of the dialogue; what AI is; and some examples of how AI could be used in local government[2];
- Responded to five case studies (see chapter 6), fictionalised examples of uses of AI in local government such as note taking in social care settings; efficiencies in planning; supporting the delivery of customer contact centres; for transport planning and infrastructure; and for identifying those at risk of homelessness;
- Noted down their daily interactions with local government services, and then considered how AI could be used in their delivery;
- Had an AI prompted conversation on their hopes and concerns for AI in local government.

Participants were also invited to submit questions about AI which were shared with the whole group. ai@cam responded to some of the more technical questions, others were discussed by participants amongst themselves.


[Figure 1: Phase 1 programme overview]

[1]HVM used Recollective, a qualitative research platform, on which HVM develops bespoke activities for our dialogue participants.
[2]Jess Montgomery, Director of ai@cam, gave an introduction to the purpose of the dialogue; Neil Lawrence, Chair ai@cam and DeepMind Professor of Machine Learning at the University of Cambridge, gave an introduction to AI, and Jess and Neil, had a conversation about topics to consider in relation to AI in local government.

Participants were also invited to submit questions about AI which were shared with the whole group. ai@cam responded to some of the more technical questions, others were discussed by participants amongst themselves.

A sub-set of 29 participants (broadly reflective of the UK population) were invited to attend a workshop to explore the discussion points raised in the online space in more depth. The group met in person in Cambridge two days after the online discussions ended. Most of the workshop involved participants exploring key topics in small group discussions. To set the scene Jess Montgomery, ai@cam's Director, provided an overview of why the dialogue had been commissioned. Professor Jennifer Schooling, Professor of Digital Innovation and Smart Places at Anglia Ruskin University shared some examples of AI use in local government, and gave her perspective on the pros and cons of these approaches. In discussions facilitated by HVM, the group considered the key dialogue questions and explored in detail what AI in local government, aligned with public needs, values and expectations, should look like.

In March 2026 a second cohort combining participants in the online deliberations, including some of those who took part in the in-person workshop on 14 November 2025, will attend a second in-person workshop in Cambridge. A final report will be produced by HVM on the whole process in spring 2026. It will set out in depth participants' public needs, aspirations, values and principles.

This interim report summarises the main headlines from this first phase of the programme.

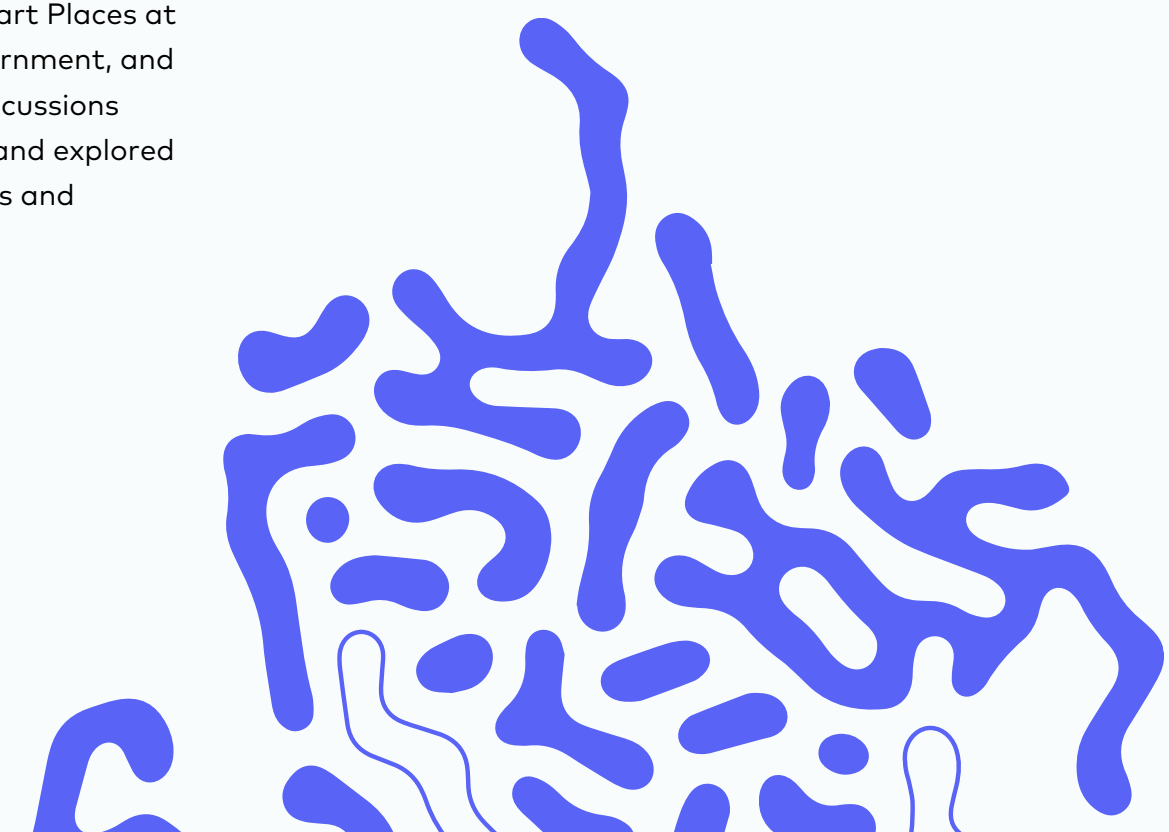

ai@cam

# What the public wants from AI in local government

Participants are united in their view that AI in local government should lead to tangible improvements in quality of life, not only cost savings or private profits. They see great potential in AI saving councils time and money which can be reinvested to make much needed improvements to services. The starting point for decision making about where and how to deploy AI should be based on where needs and benefits are greatest, not on where cost savings are most likely, or on where the technology can be used in the most novel or profitable ways.

## AI in public services should enhance what humans do, not replace them.

Participants feel strongly that human oversight of AI systems is essential (see chapter 5). They believe councils can achieve better outcomes if AI does what machines are good at, for example routine, repetitive tasks, and that this frees up staff to do what humans do best; showing empathy, creative thinking or sensitive decision making. Participants hope AI will improve job satisfaction and staff retention, as employees in an AI enabled service would spend less time on paperwork or bureaucracy, and more time using their skills and expertise. They also hope AI will lead to better human interactions in public services, as opposed to an overreliance on technology, which would exacerbate social isolation.

## Effective, user-friendly and reliable services should be ensured from the outset.

Many participants expressed frustration with existing AI applications, particularly chatbots, which they feel are a cost-saving measure, and are not yet sophisticated enough to answer even routine queries effectively (see case study 3 on contact centres). Participants talked about the importance of developing and deploying AI systems only when they have been proven to provide a good level of service.

## Inclusion, choice and fairness should underpin the design of all services where AI is used.

Participants expressed concern for less IT literate members of the community; those living in digital poverty, or vulnerable, elderly, unwell or disabled people who may struggle with new AI systems. They agreed that the option for human interaction with the council should be easily available. Some participants were excited about the ways AI could make life easier for neurodivergent staff or members of the public, and those whose first language is not English (See case study 1 on social care). Participants hope that uses of AI in public services can ensure fairness by striking a balance between standardising systems, and enabling staff to make nuanced decisions using human judgement (see case study 2 on planning).

## Joined-up and updated systems should provide better user experiences and quicker responses to residents' needs.

Many participants expressed the hope that the introduction of AI would not simply replicate existing inefficiencies and bureaucracy, but would lead to new ways of doing things, and a proactive rather than reactive approach (see case study 2 on planning). Many participants would like careful data sharing across services to streamline interactions with the council and make services more consistent. They hope that knowledge sharing and collaboration across local authorities could bring new ideas and achieve economies of scale.

ai@cam

# What the public expects from local authorities using AI

## Participants expect local authorities to involve service users and the general public in the development of AI systems.

Many would like to see a diverse cross section of local service users and the general public involved at every stage, from the development of guidelines and principles, to system design, testing, monitoring and ongoing improvement.

## Jobs and livelihoods should be protected.

Participants are very concerned about AI in public services leading to job losses, and some would like to see concrete measures put in place to protect jobs in the coming years. Many see the potential for AI systems leading to job creation in roles such as editing, monitoring inputs and outputs, or cyber security.

## Council staff should receive quality ongoing training to perform their new roles well.

Participants feel strongly about the importance of human oversight in roles where AI is used in public services. This stage of their role should be ring-fenced, valued and given time, to prevent over-reliance on AI, loss of skills and capabilities, or self-fulfilling prophecies in decision making where AI has made suggestions. In this new context, staff should be trained in new skills such as fact checking, decision making using AI data and actively bringing in human qualities which will be missing from what AI provides.

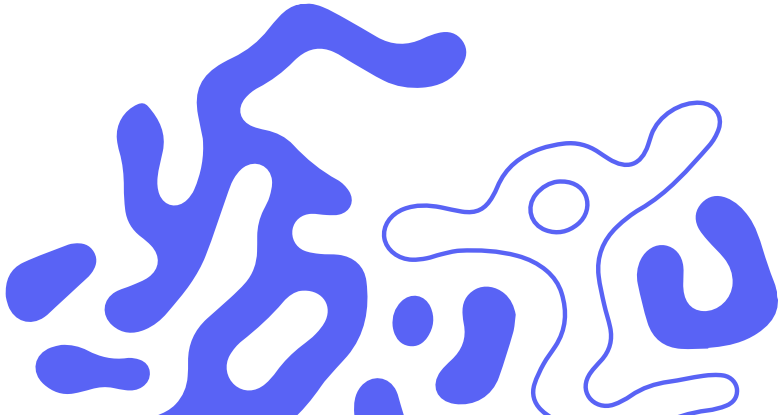

## Local authorities and developers should be responsible and accountable.

Participants would like to see enforceable regulations and clear guidelines in place to ensure services are high quality, ethical and secure, and to prevent misuse, mission creep, corruption or profiteering. Participants expect AI developers to work directly with local authorities, rather than councils buying in systems developed by unaccountable, profit driven multinational corporations. Importantly, participants call for a named person to be clearly responsible for every application of AI, so mistakes can never simply be blamed on technology.

## Building and maintaining trust is essential at multiple levels.

Participants identified trust as a critical foundation that operates across several dimensions. They need to trust in the AI systems themselves - that they are accurate, reliable and won't be manipulated to produce fake or misleading outputs. They need to trust in local government - that AI won't be used for inappropriate surveillance, that contracting processes will be free from corruption, and that their councils will act in residents' best interests. They need to trust in the companies or organisations developing AI systems - particularly around their motivations, ethics and accountability. Finally, participants call for independent regulation that they can trust to enforce standards and protect public interests, rather than relying solely on local authorities or developers to self-regulate.

## AI systems should be planned, tested and monitored robustly; improvements should be ongoing.

Participants emphasise the importance of long-term planning and stringent testing before any new AI systems are rolled out. They advocate for step-by-step introductions, alongside ongoing independent monitoring and evaluation. They talk about the need for a flexible and agile approach to make constant improvements in this new, fast-paced and ever-changing field. Participants stressed the principle of 'hope for the best but plan for the worst' - while optimistic about AI's potential, they feel it is essential to have clear exit strategies and contingency plans if AI systems are causing unforeseen problems, or not working as they should be. This includes maintaining staff capabilities to continue service delivery without AI if systems fail or are compromised. Participants felt that AI should be introduced incrementally in local government services. They want to see systems proven to work well in lower-risk applications before being scaled up or applied to more sensitive areas. This measured approach would allow for learning, adjustment and building public confidence before expanding AI use across services.

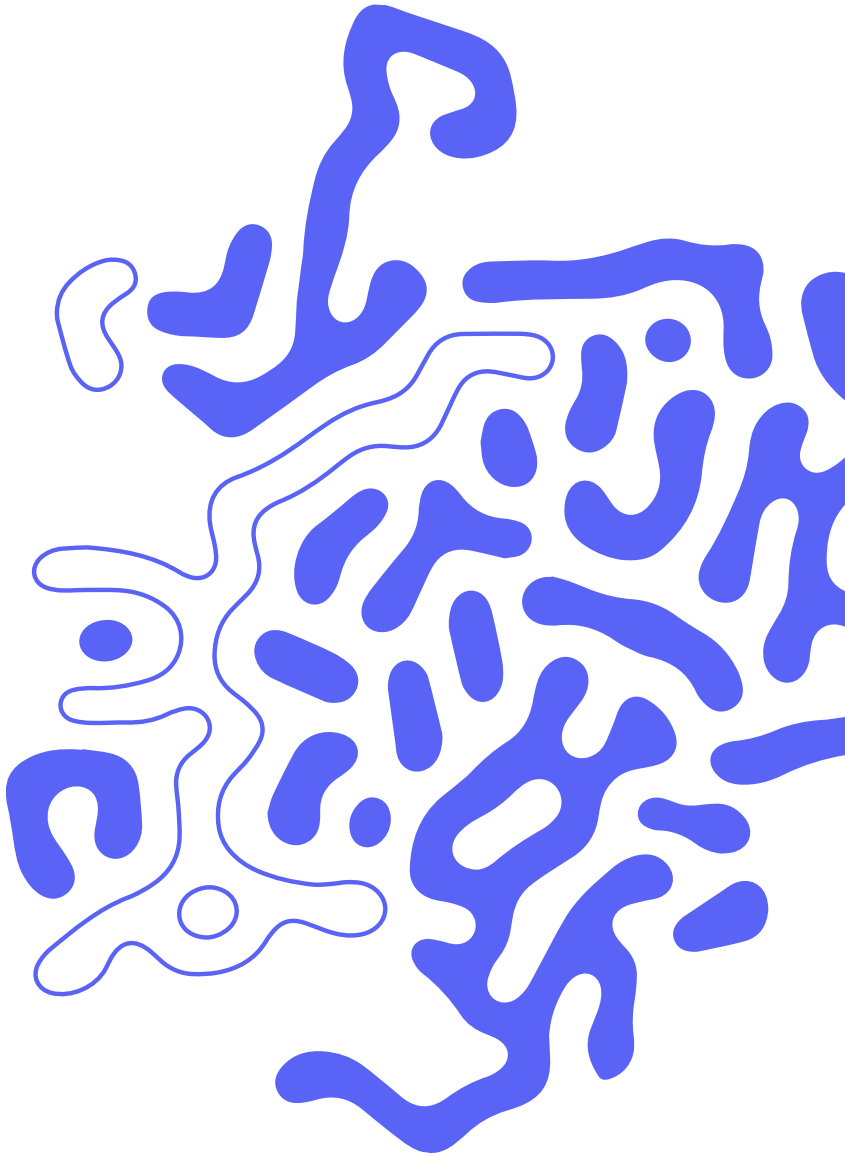

## Transparency is essential at every level.

Participants feel strongly that the public should always be informed when AI is being used in their interactions with the council. They also need clarity around the ways their data is being used and stored. Many participants point out that comprehensive information about the behind-the-scenes workings of AI in local government should be in the public domain. This should include costs in terms of both the running of AI, how these costs are being met, and, importantly, where any savings are being made and reinvested, as well as details of who the developers are and how they work.

## Local authorities should strive for accuracy and impartiality.

Participants point out that the quality of information going into AI systems should be carefully monitored, and sources should be traceable. Mechanisms to identify and rectify mistakes should be in place, and several participants suggested that local authority systems should be trained to pass users on to a human being if they cannot answer a query, rather than potentially 'hallucinating' incorrect responses. Participants say AI systems should learn about and reflect the complexity and diversity of human communication, and be trained by diverse teams of people to prevent bias. This should include political bias; AI should not reflect the political allegiances of any local council.

## Security and data protection should be well resourced and privacy issues anticipated.

Many participants expressed concerns about whether cash-strapped councils would be able to afford the world-leading security and data protection systems necessary for the highly sensitive data they are likely to gather. With the introduction of AI, participants point out that this not only applies to personal details, but many other types of data such as behaviour patterns (see case study 4 on transport and infrastructure). They are also concerned about the safety implications of malicious actors hacking into large-scale networks such as traffic systems. Participants suggest that in addition to robust technology-based backups, council staff should keep the skills and capabilities to do their jobs without AI in the event of systems being compromised or malfunctioning. They say the public should also be protected from the increased possibility of scams and cybercrime if AI becomes part of interactions with their council.

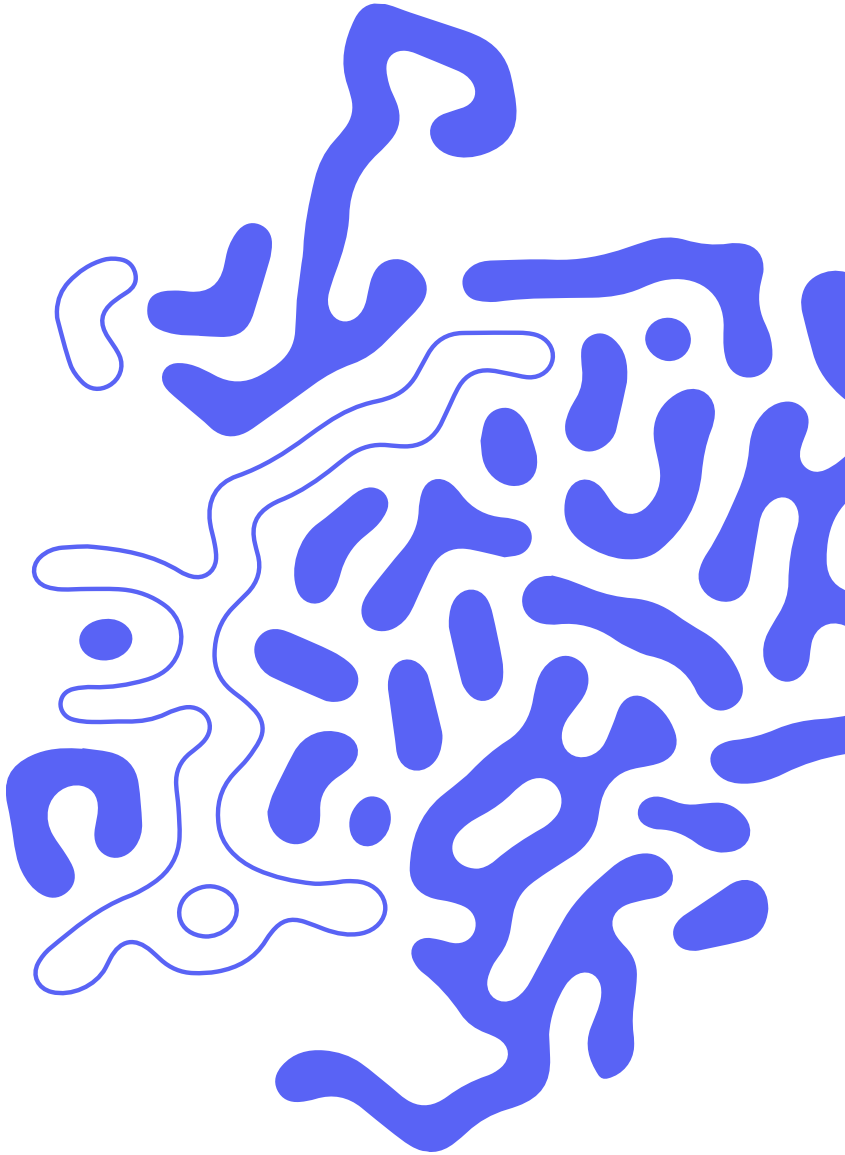

## The climate impacts of AI in public services should be measured, communicated and mitigated.

Participants are aware that data centres and other infrastructure necessary for AI are highly energy intensive. In the context of local authorities ambitious Net Zero goals, participants call for transparency over the climate impacts of AI, as well as measures to mitigate these impacts, including being selective about where and how AI is deployed.

## Public engagement, education and communication are essential to success.

For the rollout of AI in public services to be equitable, accessible and consensual, participants feel strongly that the public needs to be engaged and supported to understand and use new systems. Participants would also like progress with AI in public services to be quantified, celebrated and communicated, for example information on council tax bills about time or cost savings that have been made, and how they have been reinvested to improve services.

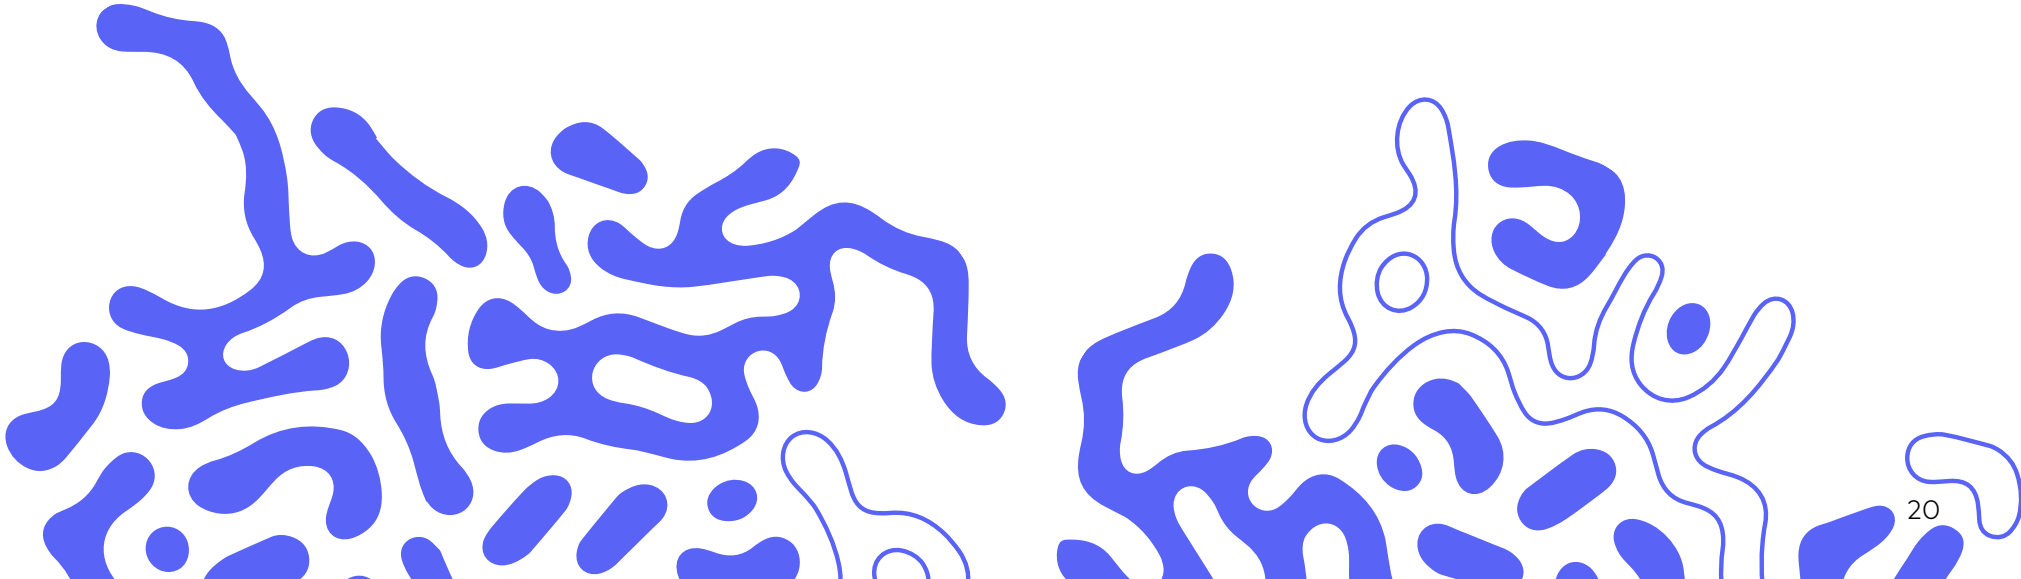


# Where and how AI should and should not be used in public services

# Case studies

On the online forum Recollective, which HVM tailored for this public dialogue, participants were asked to review five case studies showing how AI could be used by local government. They were asked to share their feelings, hopes and concerns about each one, along with the reasons for their answers. The following is a summary of their reflections.

## Magic Notes in Social Care

Participants can see clear benefits in AI immediately reducing paperwork to free up social workers to focus on the people they are caring for. Several participants explained their use of AI transcription in their own work and although they say that it always needs checking, in their experience it has been very efficient. Some participants also point out the benefits of AI transcription for care workers who may struggle to take extensive notes for various reasons such as dyslexia, repetitive strain injury or English not being their first language.

Ongoing training in this new skill set is essential according to many participants. They feel that careful checking, editing and human decision-making are a non-negotiable part of using Magic Notes. This part of the role should be valued and given time, to prevent over-reliance on AI which could lead to error, or unnecessary personal information being recorded and stored.

Body language, nuance and tone cannot be picked up by AI, but may be important to record in a care assessment, so some participants feel care workers should continue to take notes. There may also be situations where using Magic Notes is not possible due to systems being down or recording not being appropriate, so care workers should not lose this skill.


### The challenge

Jo works in social care, where she spends more time filling in forms than she does caring for people. Just one assessment visit in someone's home generates several days' worth of paperwork.

Valuable colleagues who were empathetic and compassionate, left the job due to the pressure of the admin workload. It was especially tough for colleagues with dyslexia.

### An AI application

Jo's local authority just introduced Magic Notes[1], an AI tool which instantly transcribes and summarises the conversations she has during an assessment. This frees her up to focus properly on the person in front of her. She no longer needs to write up lengthy case notes, and many of the follow-up forms get populated automatically. Magic Notes also suggests or takes care of actions like drafting letters to GPs.

Jo takes time to review the summaries and the actions proposed, as these documents can be potentially life-changing, and they contain a great deal of sensitive information.

[1] Magic Notes has been developed by Beam, a company which develops tech-enabled welfare services for local government

*This case study is inspired by real AI applications but is not directly attributable to any council or individual.*

[Figure 2: Case study - Magic Notes in social care]

## AI to streamline planning

AI in planning could bring more consistency and free up officers to better use their expertise. Participants can see the benefit of AI taking on administrative preparations, and providing access to databases of past applications to help make decisions more consistent. They say planning officers could then bring their understanding of community contexts, emotions, styles and other specialist knowledge to the decision-making process.

Many participants expressed the belief that the planning system is outdated and overly complex, and AI should not simply replicate this bureaucracy. They know from experience that there is a great need for more efficiency and consistency, but feel this needs to come from simplifying the system for both applicants and for the council. Some participants say the complexity means applicants use AI to understand planning, find justifications for their cases, automate appeals and even play the system, which could increase workloads and lead to unforeseen issues.

Participants are also concerned that minority voices and detail could be lost in AI summaries of large amounts of comments and objections. Systems would need to be designed and officers trained to mitigate this, with officers having access to source information to weigh up who is making objections and why, especially in cases where vested interests or power imbalances are part of the context.

Case study:
AI to streamline planning

### The challenge

Chen is a planning officer who doesn't have enough hours in his day. Every local construction project, whether it's replacing a window or building an extension, requires the submission of complicated forms and technical drawings. He needs to manually validate each of them, while his backlog of applications grows, and his inbox fills with requests for updates. He gets frustrated having to reject applications due to the same simple mistakes; incomplete sections of forms, missing floorplans, or incorrect application fees.

### An AI application

His local authority are introducing an AI system designed to speed up this evaluation process. It can digitise and read different data types, like drawings, text or spreadsheets. This means it can detect and even correct common mistakes, generate summaries, and create searchable databases of plans. It provides references, linking summaries back to original submissions, so Chen will have oversight of all decisions. It can also search through past decisions on similar applications, to help ensure more consistent decisions in the future.

[1] The University of Liverpool is working on a project to develop AI systems for this purpose, and to facilitate the consultation process on planning applications.

*This case study is inspired by real AI applications but is not directly attributable to any council or individual.*

[Figure 3: Case study – AI to streamline planning]

## AI-powered customer contact centres

Many participants feel that the purpose of AI attempting to emulate human contact in this way is to cut costs, not improve services. Participants shared a range of frustrating experiences with chatbots and AI voices agents failing to answer their questions, giving false information, misunderstanding regional accents, taking conversations round in circles and wasting their time. A few had positive experiences, and suggested chatbots can be one of many tools available to find information. But the dominant view is that customer-facing AI is not effective enough yet to be the first port of call for residents, and can be a barrier to residents getting the help they need.

Participants point out that if someone is phoning the council, it is likely that they need to speak directly to a human. Several participants explained that they only us council phone lines when they have exhausted other avenues, such as the council website, google, chatGPT or YouTube. Many pointed out that vulnerable or elderly people may be the ones more likely to use a phone line, and that they may need the care, understanding and familiarity of a human voice.

If routine queries cannot be answered on the website, the website should be improved first. Some participants talked about the difficulty of navigating council websites to find the answers to routine queries. Several suggested an advanced AI search feature on the website would be useful if it could draw on all non-confidential council documents, such as meeting minutes, and answer both routine and more obscure search queries.


[Figure 3: Case study – AI to streamline planning]

## AI for transport and infrastructure

Participants are enthusiastic about these uses of AI: synthesising different data types in a way humans cannot, to solve genuine problems. In the case of traffic management, participants can see great potential in reducing congestion, emissions and accidents by optimising routes for cars and buses, using real-time data on traffic volumes, bus demand, roadworks, weather or local events. In the case of potholes, sources including images, traffic data, weather patterns and resident's reports can be synthesised to detect potholes early. Participants in all small groups felt this would be a particularly valuable use of AI.

Pothole identification is seen as a great use of AI because it is perceived to have no emotional or moral component. But it must lead to action. Participants can see the potential of AI taking care of reporting and classification to allow staff to focus on solutions. They also appreciate the efficiency of repurposing existing services like bin collections, as long as they do not amplify existing inequalities in rural-urban service provision. Participants emphasise that investment in AI technology is of no use without investment in staff and materials to ultimately fix more potholes.

Monitoring behaviour brings mixed reactions: smart traffic lights and preventing poor driving are desirable, but people question whether this could compromise privacy and if the data could be misused. Many participants are uneasy about an increased use of cameras in public spaces, without the public's explicit consent. Participants call for clear boundaries for what this data will and will not be used for. There are also concerns that recording patterns of behaviour, such as when people are not at home, could compromise their safety.

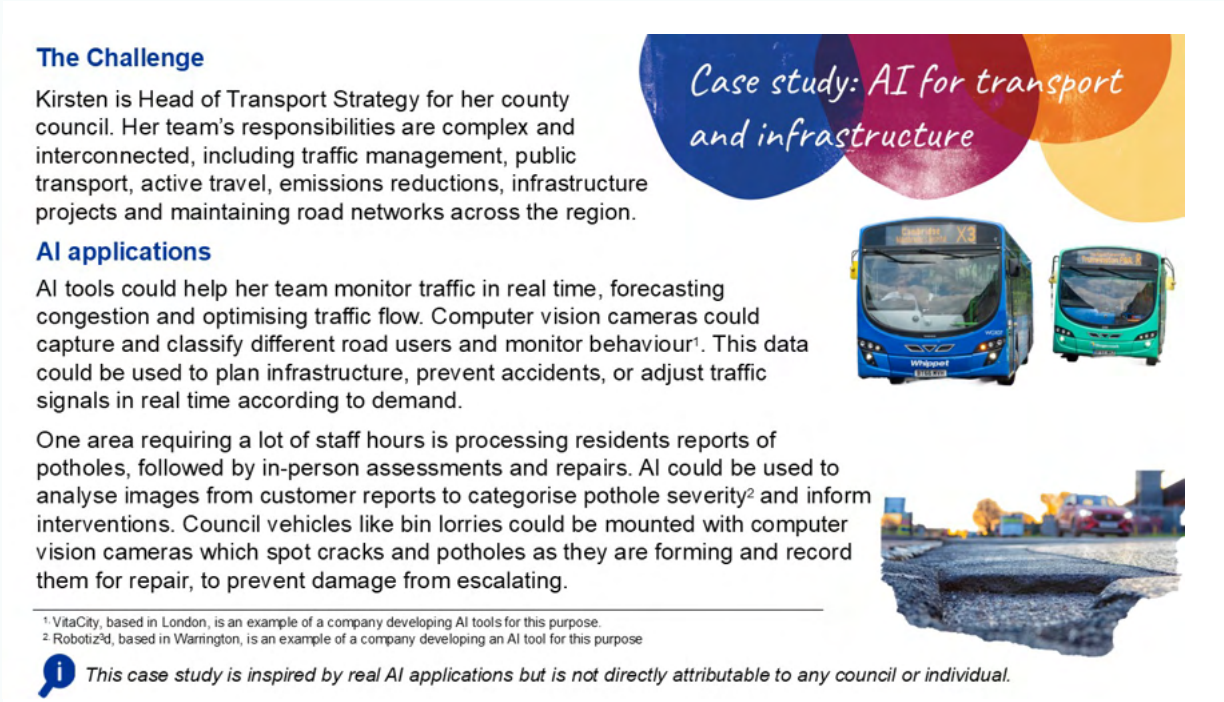


[Figure 4: Case study – AI for transport and infrastructure]

## Preventing Homelessness

Reactions to proactive support in this context are mixed. Several participants emphasise the importance of preventing homelessness because of the lasting trauma and damage it can cause. Others raise the possibility that this kind of intervention may make vulnerable people feel targeted or humiliated.

Many participants find this use of AI problematic due the vulnerability of the people involved and the complexity of their situations. They point out that people at risk of homelessness who have left prison or have mental health, substance abuse or literacy issues, may find it difficult to be monitored in this way and may not have actively consented.

Participants feel that many characteristics may lead to discrimination, beyond those that are protected. Some argue that protected characteristics could help to paint a more holistic picture of a person's risk of homelessness and that other less obvious characteristics may lead to just as much discrimination.

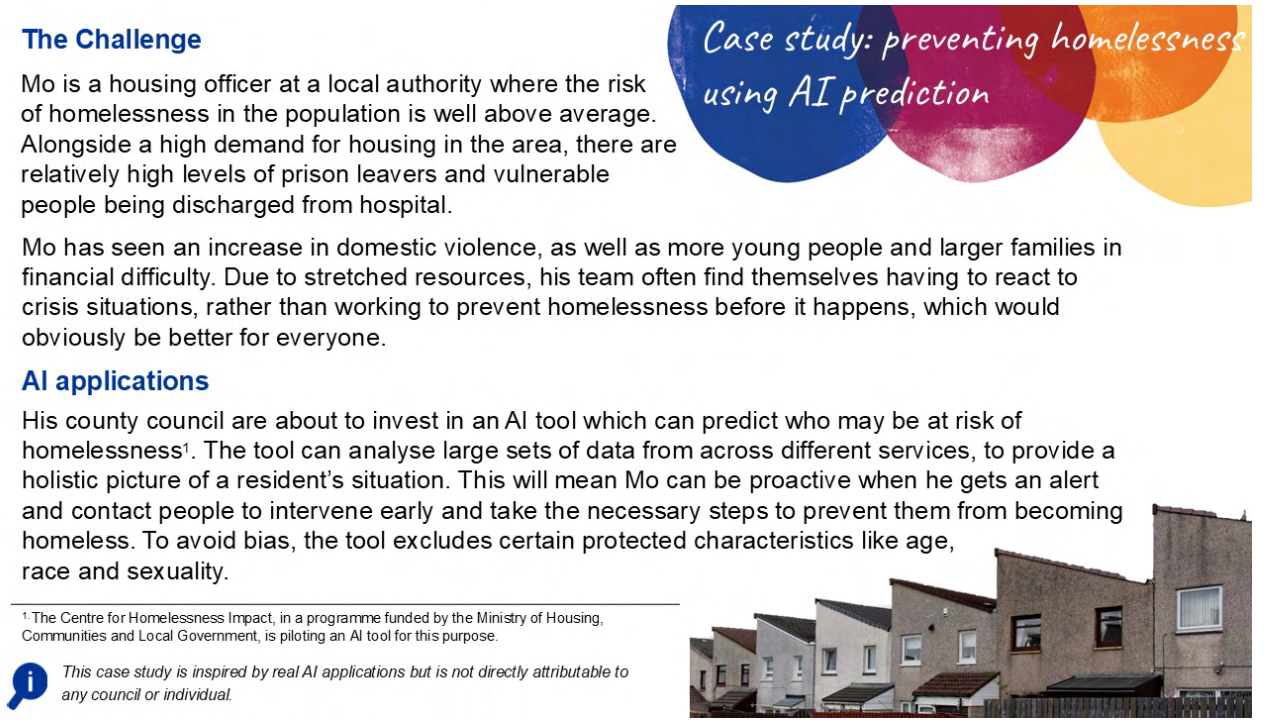


[Figure 5: Case study – preventing homelessness using AI prediction]

ai@cam


# How AI could be used in other areas of local government

ai@cam

As part of their online Recollective space before the in-person workshop, participants were asked to keep a 5-day journal of their interactions with local government services and think about where AI could have helped (or not). In the workshops, participants discussed this further, thinking about how AI could be used in local government to best meet the needs, values and expectations of residents and why AI might be useful in those cases.

Their ideas on the nature of AI action are summarised below.

## INFORM AND ANSWER

**Providing information in response to resident queries**

- A chatbot or an AI-powered search feature answering questions
    - Including about council-run events like a fireworks display or music events
    - Questions about council tax banding or exemptions, roads not maintained by the council, which bins to put items in, live bin collection routes
- Transport monitoring and journey planning e.g. accurate bus waiting times
- Providing information about upcoming local events, services and businesses in an area
- Finding books in libraries across the region
- Providing information about busy times at leisure centres, or number of spaces left at Children's Centre activities

## TRANSLATE

**Translate between languages**

- Translating documents e.g. making local government materials accessible for Ukrainian refugees
- Making Simple English or easy read versions of local authority documents

## RECORD

**Making records or reports**

- Note-taking from a recording of a meeting, client visit or planning committee
- Allowing residents to report issues like potholes, or send requests for new parking controls
- More interactive application processes e.g. to housing register
- Processing applications for parking permits, blue badges or bus passes
- Processing complaints

## SIMULATE

**Modelling the effect of changes**

- Simulating the impact of development for planning consultations
- Simulating potential effects of new road or town centre layouts on local businesses, traffic flows etc
- Simulating new services

## SUMMARISE

**Summarising and counting data**

- Processing responses to public consultations or planning applications
- Counting local election votes

## AUTOMATE/PREDICT:

**Supporting service delivery**

- Waste collection e.g. with sensors in bins to monitor how full they are to detect missed collections, or the most efficient routes for bin lorries to take
- Transport management e.g. managing traffic flows, detecting faulty lights
- Street or highway lighting e.g. based on most likely routes taken by a pedestrian or motorist
- Predicting frequency of maintenance tasks
- Automating local-authority managed car parking payment process, or car parking space allocation
- Predicting anti-social behaviour 'hotspots' to allocate resources effectively
- Allocation of school places
- Clean data e.g. find duplicates, identify missing data or errors

# Participants feel AI could be useful in these cases, based on five main factors:

**Providing a quality service:** Participants feel that if AI were to take some of the routine administrative burden of local authority staff's work, then staff would have more time to focus on delivering quality services and in their interactions with residents. Beyond the obvious benefits of a better service, participants feel this could enable a better and more trusting relationship between local authorities and residents.

**Speeding things up:** Participants feel that AI could greatly speed up the time taken for some tasks. As above, they feel this could lead to improved service quality and reduced waits for services e.g. applications to be processed. Some feel that a loss of accuracy in exchange for speed could be acceptable in some situations, as long as human checks were in place.

**Making services accessible:** Participants feel that in some cases, AI could make local government services more accessible and more easily navigable for residents. The examples cited by participants included, for example, helping them to find information about services more easily, or enabling people to provide information (for making applications or requests) in a more conversational way in comparison to current systems. Participants also feel AI could make services and documents accessible to residents for whom English is not their first language through translation, allowing people to express concepts better in their first language where they need to.

**Improving people's lives:** Participants feel that if services were able to be improved through use of AI (in terms of speed, quality and accessibility) then these services might be able to more effectively have an impact on people's lives and in protecting the environment. For example, improved public transport systems might mean fewer people travelling in private vehicles and therefore reduced emissions. More dynamic systems might enable food banks to offer more fresh food, improving users' health. Clearer information about recycling could improve rates of recycling and reduce resource use. Quicker planning processes could enable house building to reduce the pressure on housing.

**Making best use of resources:** Whilst participants are clear that they do not want cost savings to result in less spending on services, or staff to be cut, they do feel that AI could be used to cut waste and duplication and make best use of resources. This is especially the case where participants feel that AI could be used to join up services or information.

# Grey areas and red lines: where there is public concern or disagreement and where AI should never be used

# Decision-making

In general, there were few hard red lines expressed by participants. However, many participants feel that AI should never be used to make final decisions. They feel that AI should be a tool which humans used, which may even be able to recommend a course of action, but any final decisions should be taken by humans.

However, when explored further, some groups nuanced this by saying that AI should never be used to make final decisions about people or things which directly affected their lives, and especially where vulnerable children and adults are involved. Participants feel particularly uneasy about AI making final decisions about financial matters. Other examples that participants gave included school exclusion appeals. In contrast, they could imagine circumstances in which they feel it would be appropriate for AI to make decisions, when these were about infrastructure which only indirectly affected people's lives e.g. an AI tool grading the severity of pot-holes and making a decision to send a team out to fix. These uses feel more like automation of routine processes, with fewer consequences for people's lives.

While some participants feel there were clear areas of council services where AI use was not appropriate (e.g. housing, social services, legal matters, health and benefits[3] which were all considered too complex and consequential) most feel it could be used in any area of the council's services provided it was in line with the purposes and principles described in previous chapters.

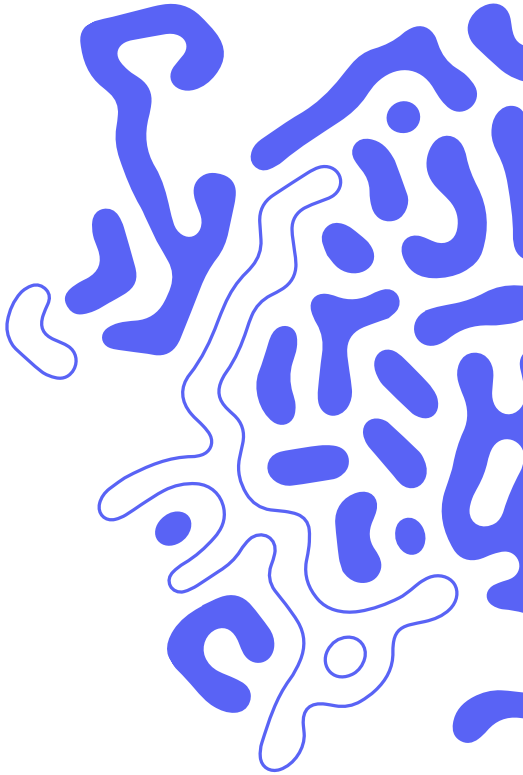

3. Participants were not at all times clear which services come under the remit of local government. For example, some raised services such as general clinical care from GP surgeries and hospitals and types of benefits (e.g. 'Jobseeker's allowance' / Universal Credit)

# Predictions

There was a lack of agreement on whether AI should be used by local government to make predictions. Again, this was not a concern when it came to infrastructure (e.g. predicting the frequency of maintenance required for a gate, or bin collections). However, many participants are concerned about uses of AI which sought to predict things about people, especially children. Participants were shown an example by specialist Professor Jennifer Schooling about AI use to predict families in need of social care support. Participants who were concerned about these uses feel that using AI tools for such purposes could turn predictions into self-fulfilling prophecies, and would not be able to take the full complexity and individuality of family situations into account. However, others saw potential in such scenarios, as being able to provide early intervention support.

Participants also expressed some red lines around how AI was used, relating to consent, purpose and security. These included:

- That AI tools should never be used without residents' knowledge and consent.
- That AI tools shouldn't use data from residents who haven't consented to their data being used.
- That use of AI shouldn't lead to confidential and identifiable information being shared.
- That technology shouldn't be trusted over people. Several participants urged those working in this space to learn from the Horizon scandal.
- That AI shouldn't be used for surveillance.
- That AI shouldn't be used where human traits (like compassion and sensitivity) are needed, or where the contact with a resident may be an important element or purpose of providing the service.

# Conclusion

This report is a headline summary of findings from the first phase of ai@cam's public dialogue on AI in local government. It was designed to help ensure the development and deployment of AI in local government services fits with public values, expectations and needs. In what quickly became a rich and lively online forum, 95 Cambridgeshire residents shared their many hopes, concerns and ideas for AI in public services over the space of 5 days. 29 of those residents then took part in an in-person workshop to discuss in more depth what is important to them in this space and why.

The findings demonstrate a cautious optimism amongst participants. Many are aware of the pressures local authorities are under in terms of budgets, staffing and paperwork and they can see huge potential for AI to save local councils time and money, which could be reinvested in improving services.

There is strong consensus that the guiding principle for introducing AI into public services must be to bring tangible improvements to people's lives, rather than to cut costs, replace people or introduce innovation for its own sake. Local authorities should be selective about where and how AI is deployed, and use the technology to do what it already does well and where it will bring the greatest benefits. For example, synthesising large datasets or routine record keeping are seen as a better starting point than uses which emulate what humans are needed and valued for in public services.

Participants want jobs in local authorities to be protected, and training and support provided to make the most of these new ways of working. They call for thorough planning, testing and monitoring of AI systems, and for transparency and accountability to underpin deployment. Accuracy and impartiality must be strived for, and security and data protection should be robust. The public should be brought along on the journey with effective communication and education.

Participants are enthusiastic about AI being used to ease the administrative burden of staff who could use their skills and aptitudes more effectively elsewhere. Areas such as transcription, translation, infrastructure image analysis and dataset synthesis seem very promising. Participants are wary about AI performing customer-facing roles, as the technology does not seem effective enough yet. Participants are clear that they do not want AI to make final decisions that directly affect people's lives, especially concerning vulnerable members of their communities.

This public dialogue continues with a workshop in Spring 2026, again with a cohort a broadly reflective of the UK population drawn from the Cambridgeshire residents who took part in Recollective activities and the November workshop, to further develop and expand on these findings with a focus on specific use cases. What is clear from both the process and the findings of this research so far, is how willing and able the public are to be involved in decision making at every stage of AI development and deployment in public services – and their concern that publics across society continue to be involved in such decisions.

# Thank you

ai@cam | UNIVERSITY OF CAMBRIDGE