


AI in Local Government

From principles to practice: insights from a public dialogue

NAVIGATION

Contents

01	● Foreword from ai@cam	03
02	● Summary	04
03	● Introduction and methods	07
04	● Public views on AI in local government and how they have evolved	09
05	● AI in local government in practice: exploring six live projects	11
06	● What the public’s response tells us	22
07	● The people’s ethical design framework for AI in local government	24
08	● Conclusion	29
09	● Acknowledgements	30



SECTION 01

Foreword from ai@cam

Artificial intelligence is reshaping the institutions that underpin everyday life. Local government is one of the most important of those institutions. It delivers services that people rely on every day, often at moments when they are most vulnerable, while operating under significant financial and operational pressure. AI offers new ways to address those challenges, but only if it is applied to the right problems, in the right way.

Universities have an important role to play in that process. They are places where new AI technologies are developed, and where technical expertise, public institutions, and citizens can come together around shared challenges.

This dialogue convened researchers, policymakers, and publics to design ai@cam's AI for Local Government Accelerator. This Accelerator supports collaborative R&D projects between University of Cambridge researchers and partners in local government that use AI to tackle live operational challenges. Its first phase helped shape the criteria we used to select projects to join the Accelerator. In its second phase, public participants considered those projects, questioning their assumptions and suggesting how to strengthen them. Public dialogue became part of the research and innovation process.

Participants wanted to understand what difference AI would make to public services, whether existing approaches could already achieve the same outcome, and how any benefits from AI in local government would be felt by residents. Essentially: is AI the right tool for the task, and will it deliver public value?

The discussions converged on a picture of what good AI in local government should look like. Participants expected AI to be used where it addresses a real public need, where it complements human expertise, where it is designed to be inclusive and accessible, and where accountability remains with people. In each conversation, participants returned to whether AI was addressing public needs, improving people's experience of services, supporting human judgement, and strengthening the institutions people rely on.

The dialogue also highlighted how public perceptions of AI are changing. For many people, the reference point for AI is now a large model developed by a global technology company with a large environmental footprint. Several projects described in this report present a different approach: task-specific AI methods, locally developed, often made open source, with data held under UK governance, and designed around clearly defined public service challenges. Making those differences visible changed how participants assessed the projects.

Across the dialogues, participants prompted researchers to reflect on the problems they were trying to solve, suggested practical improvements to project design, and encouraged us to think differently about the systems into which AI was being introduced. Their questions were often those of service designers: will this help people? Will it improve the way services work? What evidence will show that it has made a difference? They repeatedly asked us to think beyond today's projects to

whether this moment of technological change could also be used to improve the underlying systems of local government.

As AI technologies advance, universities have a responsibility to connect innovation to the needs, priorities and values of our communities. We hope this report contributes to a wider discussion about how universities, public institutions, and citizens can work together to shape innovation in the public interest.

Thank you to everyone who took part in this dialogue. Participants gave their time generously, engaged thoughtfully with complex questions, and were willing to engage in discussions with both researchers and local authorities. This report—and the projects it has helped shape—are stronger because of their contribution.



Jessica Montgomery,
Director, ai@cam



Neil Lawrence,
DeepMind Professor
of Machine Learning,
University of Cambridge

SECTION 02

Summary

Residents from across Cambridgeshire came together in a public dialogue between November 2025 and May 2026. They deliberated on the development and deployment of AI in local government services, with the goal of connecting policy and action with public values, expectations and needs.

ai@cam, the University of Cambridge's flagship mission on AI, commissioned specialist research agency Hopkins Van Mil, to deliver this public dialogue in two phases. The first engaged 95 Cambridgeshire residents on a dedicated online platform, to explore use cases for AI in local government, and express their hopes and concerns. 29 of these residents then came together in an in-person workshop in November, to deliberate in greater depth and develop a set of principles and conditions for AI in local government.

The findings of this first phase informed the selection criteria for funding and support for a new AI for Local Government Accelerator programme being set up by ai@cam. The Accelerator supports collaborative projects between University researchers and local government to deploy AI in support of local services.

Six projects were selected to join the Accelerator, and became the focus of the second phase of the dialogue. 23 participants from the first workshop rejoined for the second phase, along with 3 participants who had taken part in the online space only, selected to boost the sample of participants from lower socio-economic groups, particularly those with experience of social housing.

The 26 participants took part in online discussions before coming together with the project leads, for an all-day workshop in May, to discuss where and how these live projects are responding to the public's values, expectations and needs, and to inform their design and development. These real-world examples fed into participants' development of an ethical framework for AI in local government.

Described in this report are a set of findings on what participants are optimistic about, and where they think there should be care taken due to the potential risks.

95

Cambridgeshire residents on a dedicated online platform

29

Of these residents then came together in an in-person workshop in November

6

Projects funded through the Accelerator

26

Participants took part in online discussions before coming together with the project leads, for an all-day workshop in May

SECTION 02

Key findings: Public views on AI in Local Government

There is cautious optimism about how AI in local government can make public services more efficient, and a hunger for tangible improvements in people's everyday lives. However, the potential environmental, social and economic trade-offs of AI deployment call for a selective approach, with planning and evidence to show how wider impacts will be mitigated.

Where and why AI should be used

AI should be used where real need exists, and where it has the capacity to help solve issues that matter to the public in a tangible, measurable way.

AI should be saved for tasks that machines do well, enhancing what humans do, not replacing them.

Protecting people and jobs

Efficiency gains should be reinvested to improve services and working conditions, including with job creation and retraining. Every AI application should include clear strategies to mitigate impacts on current and future human resources.

Inclusion, choice, and trust

Inclusion and fairness should underpin all AI design and deployment, with bias actively rooted out on an ongoing basis. Choice around interactions with public services is essential, especially enabling human contact. Effective, user-friendly and reliable services should be ensured from the outset.

Fixing the foundations

Opportunities for AI to improve, join up and update local government systems more generally should be taken, avoiding replicating inefficiencies or compensating for inadequate foundations.

New AI systems should aim to surface and contribute to solving wider systemic issues wherever possible.

Service users and the general public should be involved at every stage of AI design and deployment, to ensure new tools maximise public benefit.

SECTION 02

Key findings: What ethical governance of AI in local government requires

Across both phases of the dialogue, participants moved from broad hopes and concerns about AI to a more concrete account of what responsible development and deployment would require in practice. Their expectations span institutional governance, technology choices, accountability, data and operational practice:

- **Clear and durable principles for the use of AI**, for example through a Charter for AI in local government that provides continuity beyond changes in political leadership.
- **Independent oversight and public involvement**, with a broad range of stakeholders involved in shaping how AI is governed.
- **Robust planning, testing and ongoing monitoring**, so that systems can be improved as evidence and experience develop.
- **Responsible technology choices and supply chains**, with participants favouring simple, task-specific applications and placing value on local development, open-source software, lower environmental footprints and UK-based data storage.
- **End-to-end transparency**, including public information about where AI is used, how systems are designed and procured, who develops them, how data are stored, and what impacts they have. This should include active reporting of successes and failures.
- **Accuracy and traceability**, so that AI outputs and the processes that produce them can be scrutinised, with mechanisms for correction and improvement.
- **Clear human accountability and routes to redress**, with a named person responsible for each application and accessible processes for challenging decisions or seeking remedy.
- **Well-resourced data security and system resilience**, including appropriate backup arrangements.
- **Clear boundaries around data use and access**, protecting privacy and giving people confidence about how their information is used.
- **Planning for failure and change**, with honest communication and exit strategies if projects do not work as intended.

SECTION 03

Introduction and methods

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.

ai@cam is the University of Cambridge's flagship mission on AI for science, citizens, and society. The University of Cambridge is home to a dynamic AI community pursuing high-quality interdisciplinary collaborations that take AI from the lab to areas of need. ai@cam is building an infrastructure for interdisciplinary AI innovation through research, education, and innovation programmes that leverage the world-leading research taking place across the University and connect it 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 twenty-one 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 leads to actionable insights, collaborative solutions and evidence-based policy making.

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

- Help ensure research and 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.

A recruitment specification was devised to recruit ~100 people. 95 people from Cambridgeshire, broadly reflective of the UK population joined the first phase of the process and were supported to take part in activities and discussions in a tailored online space¹. 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 government²;
- Responded to five case studies; 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.

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 Jessica 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.

The findings from this first phase were used to inform the selection criteria for funding and support from the AI for Local Government Accelerator programme. Six projects were selected for funding and support, which became the focus of the second phase of the dialogue:

- MAPLE: Map Automation for Planning and Local Efficiency
- PRISM: Predictive Risk Intelligence for Social Housing Maintenance
- SHIFT: AI-enabled Surveys for Housing Intelligence & Future Trajectories
- Deep Learning for Fly Tipped Waste Detection
- Bidding Behaviour and Outcomes in Choice-Based Lettings

¹ HVM used Recollective, a qualitative research platform, on which HVM develops bespoke activities for our dialogue participants.

² Jessica Montgomery, Director of ai@cam, gave an introduction to the purpose of the dialogue; Prof. Neil Lawrence, Chair of 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.

→ Human-Oriented AI: Design Framework for Reaching Vulnerable Tenants

In May 2026, 26 participants from the first phase of the dialogue were reconvened as part of an iterative process to discuss where and how these live projects are responding to the public's values, expectations and needs, and inform their design and development.

The second phase began with participants rejoining the online platform, where they watched short introductory videos from each project lead explaining their work, and shared their first reflections. This process included a second AI prompted conversation about their hopes and concerns for AI in local government.

The following week, on 15 May, participants and project leads came together for an all-day workshop. Participants began the workshop in small groups facilitated by HVM, to discuss how their views on AI in local government had evolved since the first phase and share first reflections on the projects they were introduced to on the online space.

Project leads then presented their projects verbally to participants and answered questions. Participants went on to explore each one in poster sessions, where they could discuss their hopes, concerns and recommendations directly with the project leads. HVM produced explainers of each project for the poster session, which were used alongside visuals produced by the project leads.

These real-world examples fed into the subsequent small group deliberations where participants built on their discussions from phase one to develop an ethical framework for AI in local government, inspired by one of the projects; Human-Oriented AI: Design Framework for Reaching Vulnerable Tenants.

Throughout the dialogue process running from November 2025 to May 2026 key questions were 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? What does this look like in practice?
- 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?

The findings of the first phase were published as interim results in early 2026 and used to shape the AI for Local Government Accelerator. This report brings together findings from the first exploratory phase of the dialogue, with the second phase focussed on live projects. It shows how participants' views on AI in local government have evolved; what participants feel

AI in local government should look like in practice, and what ethical principles should underpin development and deployment. Participant quotations are reproduced verbatim and are attributed simply to 'Participant' to preserve anonymity.

SECTION 4

Public views on AI in local government and how they have evolved

How do participants see AI in local government, and what does this mean for those developing it? This section explores how public perceptions changed between the first, exploratory phase of the dialogue in November 2025 and the second phase in May 2026, when participants considered six live AI projects against the principles they had developed in phase one.

In a context of tangible resource constraints, there continues to be cautious optimism about AI in local government

Across both phases, participants shared experiences of public services failing to meet their needs, describing fragmented and inefficient systems, inadequate infrastructure, and difficulties speaking to council staff when they needed to. There is hunger for improvement, and in the first phase of the dialogue, participants discussed a broad range of potential applications of AI which could improve services, for the benefit of the public and for the local authority, albeit with clear conditions (discussed in detail in Section 7 on the public's ethical framework).

But the acceleration in the uptake of AI across society brings mixed reactions

Between November 2025 and May 2026, participants saw a rapid proliferation in AI use in their personal and professional lives. Many described astonishing efficiency gains they had experienced, using AI to draft documents, take care of administrative tasks, or help with organisation. But equally, many feel the global hype around AI is leading to overuse and a tendency, or even pressure, to apply AI to every possible task or issue, whether it is the right tool for the job or not.

“Has AI been introduced because it's useful or because somebody said you need to be using this and it hasn't actually served any purpose, it's just been thrown in?” - Participant

Participants are increasingly aware of the potential environmental, social and economic trade-offs of using AI

After phase one's wide-ranging discussion of potential applications of AI in local government, many participants returned to the dialogue with greater awareness and concern about the costs and trade-offs involved. Many spoke about the energy, water and land use associated with AI infrastructure. There continues to be concern across all groups about job losses, with an increased focus on AI replacing entry level jobs and the impact on young people. Participants also expressed concerns about AI in local government worsening social isolation. Many participants alluded to the economic and ethical implications of potential dependence on major American technology firms in UK public services.

“I didn't realise how much energy AI used, I didn't realise how much water AI took, you know the technology behind it, the infrastructure that supports it, so that's something that really has come more to my mind since we last met.” - Participant

Many participants equate AI with well-known LLMs like ChatGPT

Traits participants associate with popular large language models (LLMs) – hallucinations, high energy and water use, the scale of the systems involved, or reliance on major US firms like Google or Amazon for technology and storage – are often assumed to apply to all AI. Where this is not true of the AI being developed and used in local government, that distinction is not obvious to the public.

There is broad agreement that AI should be used where real need exists

Participants want to see a selective and discerning approach where AI is addressing the right problems which improve quality of life. Innovation for its own sake, deployment driven primarily by cost-saving opportunities, or uses that fail to take into account the root cause of an issue, will not gain public trust and support.

Many participants can see the potential for AI to save councils time and money which could lead to much needed improvements to services, but it must bring tangible benefits in the areas that matter to people most. (See further discussion in Section 5) In doing so, they would like to see the time and money gains being reinvested in areas where there will be the greatest benefit.

“I just don't want it to feel like a vanity project because the council has a habit of fixing things that aren't broken and they're not fixing the things that are broken.” - Participant

AI can be used for what machines do well, not to replace humans

In the first phase, participants were clear that AI is best suited for routine tasks, data synthesis and relieving administrative burdens, which they hope would free humans to apply empathy, judgement and creativity. Participants expressed hopes that AI could improve job satisfaction and staff retention, as employees in an AI-enabled services would spend less time on paperwork or bureaucracy, and more time using their skills and expertise.

In the second phase, some participants went further by saying that local authorities should first consider whether AI is genuinely needed before replacing tasks that are already being performed effectively by people. The broader context of impacts on employment should be accounted for and mitigated in all decision-making around using AI. Many participants expressed frustration with a lack of staff and resources in public services and were clear that AI should not become a substitute for sustained investment in the human workforce.

“If there is a desire to improve services why aren't more staff allocated? Is it down to funding or simply ignoring the problem in the hope it goes away. Why always accept there's not enough money?” - Participant

If AI is used in decisions with significant consequences or in relation to vulnerable populations, great care should be taken

Participants do not want AI making decisions that could have direct consequences on how people live their lives, especially vulnerable populations. Discussions around social housing in the second phase raised particular concerns around AI. Participants raised four distinct concerns: that AI will simplify or too narrowly categorise complex human situations; that vulnerable people may not have meaningfully consented to their data being used by AI; that automation may undermine people's agency; and that AI may further entrench biases against historically disadvantaged groups.

Inclusion and fairness should underpin all AI design and deployment

Participants are aware of the potential impacts of AI in public services on less IT literate residents; those living in digital poverty; historically disadvantaged people; or vulnerable, elderly, unwell or disabled people who may struggle with new AI systems. They expect the option for human interaction to be easily available, and for AI to never be a barrier to human contact with the council.

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. 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.

AI-enabled services should be effective, user-friendly and reliable from the outset

Many participants expressed frustration with existing AI applications, particularly chatbots, which are seen as a cost-saving measure, and are not yet sophisticated enough to answer even routine queries effectively. Participants talked about the importance of scaling AI systems only when they have been proven to provide a good level of service. There was greater optimism about effectiveness in the second phase, as participants had experienced first-hand the range of applications and speed of improvement in AI tools.

“ I'm excited for AI to progress, but with solid foundations and quality AI - Participant

AI should be used to improve, join up and update local government systems

Participants hope that the introduction of AI does not simply replicate existing inefficiencies and bureaucracy, but will lead to new ways of doing things, and a proactive rather than reactive approach to delivering public services. Careful data sharing across services could streamline interactions with the council and make services more consistent. When discussing the six live projects in phase two, participants hoped researchers and council department would avoid a siloed approach, and instead share knowledge with other researchers, developers and across local authorities to bring new ideas and achieve economies of scale.

Beyond joining systems up, AI should help tackle the wider systemic issues behind them

In phase two, participants suggested that introducing new AI systems could do more than streamline how local government works, and should contribute to new ways of tackling systemic issues in the areas where it is brought in (see Section 5).

Participants expect local authorities to involve service users and the general public at every stage

They would like to see a diverse cross section of local service users and the general public involved from the development of guidelines and principles, to system design, testing, monitoring and ongoing improvement. Many participants say they value the role this public dialogue process is playing by listening to their views, applying them to real projects, and then returning to them to check and further inform design and deployment.

“ After reading the summary report, I can see you listened to our main concerns like taking over from the job of humans, and it's going to be implemented actually listened to. It's surprising because I feel like the council often doesn't listen to us residents, so this is like a breath of fresh air. - Participant

SECTION 5

AI in local government in practice: exploring six live projects

Six collaborative projects were selected for funding and support by ai@cam as part of their AI for Local Government Accelerator. Projects were in the early stages of development at the time of the second dialogue, with the aim that participants' deliberations could meaningfully inform their design. Participants first learned about and reflected on the projects through videos produced by the project leads and posted in the dedicated online space. During the second workshop, project leads presented their projects again, and participants took part in a poster session where they could discuss their hopes, questions, concerns and recommendations directly with project leads, and these are the subject of this section.

MAPLE

Map Automation for Planning and Local Efficiency

Every year high numbers of planning documents need to be processed by Cambridge City Council and South Cambridgeshire District Council. These can include a range of images, like hand drawn, scanned or photographed maps. Currently, each map image has to be processed, digitised and uploaded manually, which is time consuming, repetitive and contributes to backlogs.

MAPLE is an AI-assisted map processing system, which can detect and isolate the map areas within submitted documents; match and align map extracts to official maps; extract and classify features added by applicants, like boundaries or access routes, and analyse text such as “proposed extension”. These functions can all be reviewed by staff, which is essential given the legal importance of these documents. Data can then be shared across planning teams in dashboards and interactive maps.

Potential benefits include a significant reduction in manual processing time to speed up application turnaround and improve service delivery; specialist staff freed up for tasks which require human judgement and decision-making, and the ability to generalise methods and models across UK councils.

HOPES AND EXPECTATIONS

Participants’ hopes and expectations for the MAPLE project are summarised as follows:

Speeding up and streamlining planning applications

Participants believe the MAPLE project addresses a genuine need. They can see how automating some of the time consuming and repetitive elements of this work could meaningfully reduce backlogs, freeing up Geographic Information Systems (GIS) Officers for more complex decision-making, and making the process faster and more efficient for applicants.

Significant, easily quantifiable efficiency gains

Participants were encouraged by the prospect of significant, measurable efficiency gains. Project leads explained that data exist on current map-processing times and estimated that the tool could reduce processing time considerably.

An appropriate use of AI

Unlike work that requires human sensitivity or an understanding of complex social circumstances, map processing is fundamentally a data related task, involving pattern recognition and alignment, which participants feel lends itself well to the use of AI.

“ This is a data project and not a human project. And that’s what machine learning is best at. - Participant

CONCERNS AND QUESTIONS

Potential job losses, particularly at entry level

Many participants are concerned about automating more basic tasks such as map processing, which may be part of entry-level roles and provide a first essential career stepping stone, either now or in the future.

“ For someone who’s just spent thousands and 3-4 years studying, they would love that kind of entry level job versus ending up back in Tesco. I think these are potential opportunities that are being taken away from humans. - Participant

Deskilling

Participants are concerned that current employees could lose the ability to process maps without the use of AI and may become more like auditors than GIS officers. This de-skilling could leave planning departments reliant on AI and lead to future GIS officers not having learned the skillset necessary to properly interrogate the AI responses.

Complacency and automation bias

Participants are concerned that the apparent efficiency and accuracy of the system could breed complacency, making it less likely that reviewers will scrutinise outputs with the care required to catch errors.

Status of AI-processed maps as official documents

Participants questioned whether AI processed maps would always carry the same weight with solicitors or other officials in the planning process.

Comparisons with LLMs and AI in general

As with other projects, participants are concerned about mistakes and hallucinations, as well as high energy and water usage of AI. Several were unsure of the specifics of the tool and asked whether ChatGPT or Gemini could perform the same task.

RECOMMENDATIONS AND FUTURE DIRECTIONS

Be clear about the advantages of MAPLE over tools based on LLMs

Many participants equate AI with LLMs like ChatGPT or Copilot because that is what they are familiar with. This brings with it concern about high energy and water footprints, as well as accuracy issues as a result of hallucination. Participants were surprised to learn that MAPLE is relatively simple to run by a small team and has a low energy footprint because it is task-specific, and that a tool like MAPLE is less likely to hallucinate because it is not based on an LLM.

Involve the public in training the tool iteratively

During the initial development phase, participants suggested sending maps back to the residents who submitted them, to verify accuracy, train the model on feedback and monitor improvements.

Leverage MAPLE to simplify the planning process more generally

Participants asked where MAPLE could further streamline planning, for example if MAPLE could reduce the volume of documents residents need to submit or consult as part of planning processes.

Reassure the public about human responsibility

It was important for participants to hear that GIS officers remain fully responsible for any outputs involving AI, and documents produced using MAPLE thereby carry the same official and legal weight. Participants were also reassured to hear how the confidence percentages are used to improve accuracy and that this system combining human and artificial intelligence can be more accurate than a system relying solely on humans.

Show what efficiency gains mean for council staff and human resources

As the project progresses, participants would like to see evidence of how it has affected relevant council staff, for example in alleviating or increasing stress by reducing backlogs or creating performance pressures, or if adequate training is provided. They would also like each project to anticipate the potential effect on employment and consider how roles could be protected or even created.

Explore how MAPLE could bring council services together

For example, could MAPLE help building control and planning departments speak to each other more efficiently by providing improved tools for them to interact with?

PRISM

Predictive Risk Intelligence for Social Housing Maintenance

Housing vulnerability is a complex and growing challenge for local authorities. It often requires early intervention across multiple services. Repairs are often reactive, so problems can get worse, affecting homes and tenants.

PRISM applies machine learning to a range of datasets to produce AI-powered Risk Hotspot Maps, which can identify at-risk buildings for proactive maintenance and targeted investment. PRISM combines data on housing condition, age, energy performance and maintenance histories; climate hazards like extreme heat and severe winter events, and anonymised social and tenant vulnerability data such as fuel poverty. The resulting Risk Hotspot Maps can help identify properties with the highest likelihood of system failure or deterioration and generate alerts for housing teams, to help guide decisions by specialised staff and prioritise resources based on the greatest need. PRISM does not operate as a public-facing platform or use identifiable tenant information.

HOPES AND EXPECTATIONS

A proactive approach to housing maintenance

PRISM could help lower maintenance costs and reduce timeframes for repair by identifying problems before they become more serious or need costly emergency intervention.

Better housing conditions and tenant wellbeing

Some participants are hopeful that over time PRISM could contribute to social housing becoming safer and healthier for tenants, improving quality of life.

“ PRISM is potentially very exciting. My past work experience with vulnerable families and individuals highlighted the need for prompt and less stressful, drawn-out applications to try and improve housing issues, e.g. mould, to ensure healthier and safer homes. - Participant

Transparency about how the council identifies and classifies properties

They say the tool must include the ability to work backwards from any decision to show how it was made, and the role AI, the data and the staff played.

CONCERNS AND QUESTIONS

The root problem may not be a lack of data

Several described lived experiences of councils knowing about social housing maintenance issues and not acting.

“ What’s the point in it when the council don’t actually care when it has been raised? My windows are the worst they’ve ever seen, my house is freezing. It’s awful. And they know there’s a problem. We’re capable of raising problems ourselves, and we do, but then nothing happens. - Participant

More knowledge without more funding and capacity to act could compound backlogs

Participants point out that information about pre-emptive repairs would be in addition to current repairs which they say councils are not managing to deal with.

“ If you struggle doing repairs now, how will you keep up with demand of doing them in advance? - Participant

“ Do you have the infrastructure in place when you have this data all collected and you’ve got your hot spots and you’ve got your repair scores for housing? Is that not going to cause a huge mountain for you to climb to get those top priorities done? - Participant

Risks in processing and combining vulnerable people’s data

Many participants feel uneasy about the way councils may define tenant vulnerability pre-emptively and use it to label individuals or their properties without their input. Participants are concerned about potentially vulnerable tenants explicitly consenting to their data being used in the ways PRISM proposes.

Fairness and transparency in prioritisation

Some participants asked how resources would be allocated where there would presumably be a large number of pre-emptive interventions, and how certain vulnerabilities may be prioritised over others.

“ I know you’re saying you don’t want to put people with respiratory problems in places that have mould, but I think there’s a point where we don’t want anyone in there? - Participant

Accounting for the environmental impact of PRISM

Participants expect the application to be energy intensive and ask whether joining up and managing existing databases could be more effective and energy efficient, rather than AI working hard to bring together disjointed systems.

RECOMMENDATIONS AND FUTURE DIRECTIONS

Provide evidence of improvements to build trust

Participants suggest this comes in several forms, including statistics on numbers of repairs before and after the introduction of PRISM, comparison of tenant satisfaction surveys, and other qualitative research.

Involve tenants in the development and potential rollout of PRISM

Several participants emphasised the importance of working on

an ongoing basis with tenants who can ensure services are fair, transparent and designed around their needs.

Expand to rural areas and even public infrastructure

Some participants feel rural areas are frequently overlooked in policy and infrastructure investment. They also see potential in expanding the model beyond social housing to public buildings such as hospitals or schools, where predictive maintenance could reduce the risk of costly and disruptive failures.

Give tenants access to a property's repair history

Participants see this as a tool that could be used to give potential tenants a transparent look at the building history to avoid unpleasant surprises.

SHIFT

AI-enabled Surveys for Housing Intelligence & Future Trajectories

Local authorities are required to prepare an annual Housing Trajectory, which is a critical evidence base for planning decisions. It involves preparation, distribution, follow-up, and analysis of around 150 site-specific questionnaires sent annually to developers, agents, and housebuilders. While robust, the current process is manual and repetitive, needing significant staff time in administrative coordination rather than professional judgement, risk assessment, and strategic planning.

SHIFT aims to develop a human-in-the-loop AI tool that helps draft, create, and distribute surveys and questionnaires, and analyse complex structured responses by generating insights while keeping human expertise central. Potential benefits include faster more robust housing delivery evidence; reduced administrative burden on planning officers; improved experience for housebuilders, developers and agents responding to surveys, and more consistent, auditable, and reusable housing delivery data

The approach is intentionally transferable to housing trajectories that are required nationwide. Clear documentation and open-source components will support adoption by other authorities.

HOPES AND EXPECTATIONS

A 'straightforward' way of improving efficiency

Some participants say that typing up surveys manually seems old fashioned, and the efficiency gains of using AI to compile the housing trajectory make sense and are likely to save time and lower costs for councils. The kind of responses involved - housebuilding related data submitted in a professional context - is not sensitive, lending itself well to AI-enabled analysis.

“ Sounds like a really sturdy use of AI and great support for councils. - Participant

National scaling could bring significant benefits

As the Housing Trajectory is mandatory across all local authorities, this project could be easily replicated and bring significant efficiency gains at a national level.

“ The interesting aspect here is that this is a problem that affects every council and every county in the country. - Participant

Iteration built into the design

As well as council staff checking AI outputs, the opportunity to check and refine or adjust AI summaries or analyses with survey participants is an important step in the design.

“ Have it kind of show you its transcript or whatever you call it. Can you go oh no I didn't mean that or yes I did and sign it off? - Participant

Accessibility and ease of use

Participants expect AI-enabled surveys to have features built in that make them simple, accessible and inclusive for anyone, including neurodiverse or disabled participants.

Ongoing human auditing of outputs

Housing trajectories are influential documents, so participants expect ongoing independent auditing, never assuming AI can do the job on its own.

CONCERNS AND QUESTIONS

Homogenisation and missing detail

Participants are concerned AI systems analysing qualitative data, particularly those using natural language processing (NLP), could fail to represent unusual, individual or minority circumstances, in favour of broad categories and the majority view, so important details could be missed.

Environmental impact of AI

Some participants are concerned about the energy and water footprint of this type of AI, and ask whether existing software could do the same job without the environmental impact.

Impact on jobs

Participants would like to know how the efficiency gains of this project will be mitigated so that jobs are not lost or positions phased out.

RECOMMENDATIONS AND FUTURE DIRECTIONS

Expand to other sectors

There may be similar types of qualitative surveys in other areas which do not deal with sensitive data, where the R&D from SHIFT could be beneficial.

Respond to changes in the housing market

One participant suggested building in features to anticipate and account for local circumstances or changing survey participants, such as major companies moving premises or demographic fluctuations.

Consider how to link to front-end issues that matter to the public

Many participants feel strongly about the lack of adequate infrastructure and services in relation to new housing developments. They asked how the changes and efficiencies achieved by SHIFT could link with wider systems in housing trajectories to contribute to tackling this issue.

DEEP LEARNING FOR FLY TIPPED WASTE DETECTION

Fly-tipping is a significant environmental, financial, and public health challenge for local authorities. Fly-tipped waste can contribute to pest problems, accessibility barriers, and environmental harm. Current detection methods rely on residents reporting incidents, which can lead to delayed waste removal and increased pressure on council resources.

Councils already collect large volumes of street-level camera footage through refuse collection vehicles (RCVs). The proposed AI-powered fly-tipping detection system uses this footage to identify incidents automatically during routine waste collection rounds. The system detects and classifies fly-tipped waste using computer vision and deep learning models before review by council staff. It automatically captures key reporting information, including location, timestamp, waste description, and photographic evidence and integrates with existing council reporting systems to support faster investigation and response. Only necessary image data is retained and processed and facial images and licence plates are anonymised through automated blurring and secure data handling practices.

If successful, the same technology could also support detection of other public realm issues such as potholes, drainage faults, street lighting failures, and anti-social behaviour.

HOPES AND EXPECTATIONS

There is great enthusiasm for using bin lorry footage to report a range of issues

Participants were particularly supportive of making better use of infrastructure that already exists; if bin lorries already have cameras attached as they make their rounds, participants suggest a wide range of issues that could be reported; potholes, broken street lights, fallen trees or overhanging branches, illegally parked cars, abandoned bikes, damaged properties, blocked drains, dog fouling, graffiti, or overflowing bins.

“ I love the bin lorries. I'm obsessed with the bin lorries. The video with the fly tipping, it was like, Oh my God, they can be used for so much stuff. - Participant

“ Look at all the things it can solve (...) Make our local environment a better place to be. - Participant

Participants see great potential in this project developing beyond fly-tipping

Many participants say that fly tipping alone is not a major issue for them. During the first phase of the dialogue, participants discussed a case study involving AI-powered bin lorries detecting potholes, which was very popular, and some are pleased to see a similar project coming to fruition, but hope it acts as a pilot for further uses of AI-powered bin lorry footage.

“ I was thinking like how big a problem is unreported fly tipping? - Participant

“ The camera is on the back and maybe it's a kind of a bit of a test to see how they can sort out the AI on something that is not very difficult or controversial and to be able to move on to bigger projects where they've now then got that base, you know, to say, right, we've got it right. We know what we're doing. We know what the public think. - Participant

CONCERNS AND QUESTIONS

Reporting may not be the bottleneck

Many participants feel that resolving and removing incidents of fly tipping is more complex and problematic than reporting it, and question whether the resources exist to tackle any increase in data on incidents.

Rural fly-tipping where bin lorries don't record

Some participants referenced news stories or their own experiences of fly tipping in areas without regular bin collections. Participants brought up concerns around fairness in relation to resources and services being available in urban but not rural areas.

“ I live in the villages, and I know that there are certain places where these fly tippers know that there's nobody nearby, there's no cameras, and they're just dumping stuff. - Participant

Privacy implications of capturing and processing data without consent

Several participants were concerned about increasing public surveillance, and local authorities capturing images of people or private properties without consent. They asked questions around who would see and process the data and what the risks could be, particularly to vulnerable people.

“ And the privacy element of how much Big Brother's watching you even more and who gets to see that? Who gets to control that? - Participant

“ If you're capturing footage, you don't know they're being captured. You've got vulnerable people like children and women running from domestic violence and it doesn't take a lot to get into a local government role. You, you don't know who's reviewing that footage and so there's an ethical issue and there's a risk protection issue. - Participant

Accuracy of AI identification

Some participants mentioned situations which could be mistaken by a computer for fly tipping, such as people moving

house, neighbours giving away furniture they no longer need, or leaving out unwanted items to be collected by the council. This raised further questions around what the effects would be on local communities if people felt uncomfortable doing these things, or if there was no longer an expectation of the public reporting fly tipping or other incidents to the council.

“ In our area a lot of people give things away on the street, so either like a sofa or a desk or a chair. And it's a really nice community-based thing. (...) I don't want to discourage that in my street. It's a nice way for people to just talk to each other. - Participant

Risk of normalising fly-tipping

A few participants suggested an element of deterrent would be necessary or people may start to assume that fly-tipping will just get dealt with.

“ People are just going to dump stuff because they know it's being cleared up. They're not going to get caught. It's not going to cost them. So there needs to be a way of tracing them back quicker as well as clearing it quicker. - Participant

RECOMMENDATIONS AND FUTURE DIRECTIONS

Expand uses and areas covered using satellite imagery and other cameras

Some participants were receptive to the suggestion by the project lead that a future direction could involve using satellite imagery for higher resolution capture of fly-tipping which could cover rural as well as urban areas. A few participants also suggested incorporating other existing camera networks such as highways vehicles or CCTV to cover wider areas than refuse collections.

Communicate about the advantages of the AI system used

In a context where many participants equate AI with LLMs, participants were reassured to hear that the software used for this application is open source and low cost, and the data is stored locally on hard drives, not in the cloud.

Or consider a simpler human role where refuse collectors use the camera networks to report fly-tipping

Several participants suggested that refuse collectors could easily use the existing camera systems to make a timestamp, geographical marker and classification when they see fly tipping or other issues. This would circumvent complications they associate with AI around identification and data collection and storage.

Consider systemic issues and causes such as expensive or inaccessible recycling or rubbish disposal

A few participants suggested that addressing the underlying causes of fly-tipping may ultimately be more effective than deploying new AI tools alone.

“ Instead of trying to implement a new system, which I imagine would cost a lot, the council could also invest in reducing the cost of collecting, at the moment it's £40, for

an additional items it's £30 pounds, so if you've got a lot of stuff to skip, it's a very high cost. If you've not got a car big enough to get it to the tip, it's tricky...So we could be investing into lowering the cost instead to make it more accessible. - Participant

BIDDING BEHAVIOUR AND OUTCOMES IN CHOICE-BASED LETTINGS

Social housing is a vital public resource, and in England it is allocated through the choice based lettings (CBL) system. Instead of the council deciding who gets which house, CBL lets eligible residents browse available properties and bid on the ones they want. Although intended to promote transparency and choice, CBL often leaves residents making repeated decisions with limited information. Outcomes depend on complex interactions between types of property, household needs, and the distribution of points, so it's not always easy for tenants to understand how the system works. Housing officers must manage fairness, expectations, and the workload created by unsuccessful bids and uneven demand.

This project proposes a platform to support transparent, fair and efficient decision-making in CBL. It will apply machine learning to nearly a decade of de-identified administrative data from Camden, to analyse demand patterns, bidding behaviour and how the points system works within the allocations system. The resulting platform aims to deepen understanding of how the CBL system works in practice. It aims to help officers understand and predict demand and behaviour, enabling more effective use of scarce housing resources. The platform also aims to improve the resident experience by making decision-making more transparent and reducing repeated unsuccessful bids.

HOPES AND EXPECTATIONS

Greater transparency for tenants in the bidding process

Participants hope this project will provide prospective tenants with a clearer picture of their chances of success, so they can make more informed choices about which properties to bid on or accept.

“ Current bidding rankings are very vague, you don't know where you stand, like 11-20 or 20 plus? - Participant

They expect the AI-enabled system itself to be transparent, so tenants have an idea of why AI may have led to particular listings or recommendations.

Design to prevent gaming the system

A few participants suggested the possibility of tenants tweaking the system for certain outcomes if they know how it works, rather than entering their genuine preferences.

Involve people with lived experience in iterative design

Participants suggest that people who are using or have been impacted by the bidding system should feed back into any new tool.

“ You should let the bidders know the reasoning why you're suggesting, because if you just give them suggestions and then they don't know why, you can say we've given you this because of X, Y and Z and then they might come round and go, well, X doesn't mean anything, so you can change it. - Participant.

Anticipate tenants' future needs

Participants expect evolving criteria such as children's ages or degenerative disease to be taken into account in the design.

CONCERNS AND QUESTIONS

Will this improve the situation for social housing tenants or lead to false hope?

Participants, particularly those with experience of social housing are quick to point out that this system does not

increase the availability of housing stock, and they are concerned it may give the illusion of choice.

“ It does frustrate me because you're focusing all this on giving people what looks like more choice. And I think that's an appeasement issue to people who are on a on a social housing list. I think that will give people false hope (...) But it's not going to make a blind bit of difference to how long people are going to have to wait for adequate housing. - Participant

Will the tool be sophisticated enough to account for complex or unique priorities?

Participants say that criteria for choice are highly subjective and may not match assumptions made by AI. They are concerned the system could categorise people narrowly and disempower tenants if AI limits their choices without enough flexibility or tenant input.

“ As a bidder, your triage of what is most important for you might be different to what AI says. If you're a single mum, you want to be close to your mum so they can help with childcare, but actually AI is looking at different parameters. - Participant

AI analysis of historical social housing data could entrench bias and disadvantage

Some participants are very concerned about the use of historical data from groups of potentially vulnerable people who may have suffered from systemic disadvantage. They emphasise the difficulties and issues with points-based social housing, and question the validity of drawing conclusions from a dataset based on a flawed and potentially biased system. For example a system which penalises rejection may lead to tenants accepting unsuitable housing, which could in turn generate poor-fit data.

“ There's so many societal biases, racial biases. How can that then not be built into that system, and repeat patterns that have been mistakes, like to class basically all people in a certain area in a certain way and then that being a massive mistake. It echoes itself. There's some patterns in that that need to be looked at. - Participant

How is this different from a filtered search on a website?

A few participants asked why AI is needed for this application, when simpler technology already exists for a comparable output.

“ I'm curious about whether or not this is the right technology for the job. If I was to go to a website and have a profile, in that profile I might say these are my filters, so I have disability, I have a child, I have houses with this many beds I need to look for. And it will serve up to me whatever I need to see. Tell me why I need AI to do something that's filtered search on the website? - Participant

RECOMMENDATIONS AND FUTURE DIRECTIONS

Involve inclusion experts in design to combat bias going forward

Participants suggest any AI analysis of historical datasets could also be analysed by sociologists and inclusion experts to help actively remove systemic bias.

Communicate around the tool using examples not just numbers

One participant suggested the benefits would be easier for people to grasp if they could hear real world case studies of how it has helped or could help people.

HUMAN-ORIENTED AI

Design Framework for Reaching Vulnerable Tenants

An AI-powered Early Warning System for Vulnerable Tenants system is being developed by the Cambridge City and South Cambridgeshire District Councils. Improving support for tenants is a priority for Cambridge City and South Cambridgeshire District Councils, but many of the systems involved do not speak to each other. Developing the technical solutions is not necessarily the 'main' challenge though. When applying AI to address this problem in the public sector, there is a need to consider what data is involved and whether it affects critical services and people's lives or undermines public trust.

This project aims to develop a shared structure for governance, ethics, public value and data, to ensure these projects benefit the people concerned, rather than being led by the technology.

This Human-Oriented ethical framework for AI is designed across four layers; governance and management; public value creation; ethics and responsible innovation, and data, digital & infrastructure. It will incorporate the views of the groups of people concerned, including tenants and council staff, so that concerns, risks, ethical trade-offs and practical implications are documented and shape the design of digital innovations.

The framework will include diagrams and descriptions of the architecture of the Early Warning System, which show how each of the four layers are addressed, for example governance flows, public value outcomes, ethical considerations, data and infrastructure specifications. This gives the technical team a shared reference for the development and evaluation of solutions which are aligned with what the public and society needs. It ensures the initiatives are governed, legitimate, and responsible from the start, so that "reaching vulnerable tenants" is done in a way that city managers, tenants, and future partners can trust.

This project inspired the development of a general ethical framework for AI in local government which is covered in detail in Section 7. This section explores elements of the Human-Oriented AI project which relate specifically to reaching vulnerable tenants.

HOPES AND EXPECTATIONS

Human empathy and social intelligence are essential to prevent harm

In any scenario where AI is used in connection with vulnerable people. Participants expect a highly cautious approach. Participants say any flagging up of vulnerability should lead to building a relationship as a first step before any action is taken.

CONCERNS AND QUESTIONS

Agency and autonomy of residents

Participants with experience of housing vulnerability are particularly concerned about systems where tenants may feel labelled as or told they are vulnerable, especially if AI is involved in making those kinds of judgements. Some said that instead systems should be in place so that they feel confident to approach local authorities themselves.

Classification of vulnerability

Several participants pointed out that vulnerability can mean many things, and it may also be defined differently across different council services. They are concerned that tenants in complex circumstances could be bluntly classified by AI without enough nuance or flexibility around individual needs.

RECOMMENDATIONS AND FUTURE DIRECTIONS

Recognise that 'vulnerable people' are not a homogeneous group, and provide choice accordingly

For example, some participants prioritise privacy and consent in their conversations around AI in local government. Others find that organisations having access to a detailed information about them is essential to meeting their specific needs. For example one participant explained how helpful it is that her bank and utilities have a personal profile for her and adapt their communication accordingly

“ Perhaps there's a personal profile, building up with their special needs. Like (my bank) recognise my number and my name and they already know that they have to speak slowly, clearly and they will take that into account. My interpretation of this is hopefully something similar the local authority can do. For me it's essential because if I say to them, I can't hear you, the next thing you know, the volumes turned up. I find it awful having to repeat myself. - Participant

SECTION 6

What the public's response tells us



What can researchers, the University and local authorities learn from participants' questions, hopes, concerns and ideas about the real-world applications of AI in these six live projects?

Focus more on everyday, resident-facing projects

Participants recognised that the six projects respond to genuine challenges faced by local authorities. They also appreciate that in the early stages of development, low stakes projects are a sensible starting point. But they would like the direction of travel in AI research and development to be responding to the public's priorities, as the need for improvement in services is great.

“ So just on an everyday basis when someone is trying to get in touch with the council and is going on the website, that is still a nightmare, so we've got very basic things that aren't being addressed. - Participant

Start with lower-risk, data-driven projects that bring public value

Although participants would like to see projects developed which improve people's daily lives, there is caution around involving AI in applications relating directly to humans and personal data, particularly vulnerable people. Projects which promise significant efficiency gains but are low risk, like MAPLE and SHIFT are seen as good starting points and these principles could also be used in more front-end applications.

Plan for impacts on current and future human resources

Participants would like to see how local government anticipates

and mitigates current and future impacts on employment, how projects may benefit current employees, and for these considerations to be applied to wider decision-making around deployment.

Show how projects are designed for inclusion

Whether the application concerns council staff and systems or the general public, participants want clear evidence that systems are being designed inclusively and actively tackling bias.

Communicate where AI tools are low-risk, low impact and local

Several of these projects use AI technologies which are developed locally, use data which is stored locally, are low cost or open source, and do not have a high energy and water footprint. These advantages are important to participants, but they are not obvious.

Show that better data and AI leads to better delivery

Participants would like to see how better data translates into public value; this applies particularly to the PRISM, Bidding Behaviour and Fly-tipping projects.

“ I think we just addressed something that applies to every single one of these, and that is, what's the problem that is being identified. Will it get resolved any faster? Otherwise, what's the point? - Participant

Be honest when projects don't work, and be willing to stop them

Participants said they understand the value of experimentation in R&D. They accept that some projects may not have the intended outcomes and participants appreciate honesty and flexibility in these circumstances.

“ Don't be afraid to talk about it, not all projects are successful and that's just life, right? If most of these are part of some sort of accelerator type process, not all of it's going to work. - Participant

Contribute to improving systems and tackling wider issues

Participants would like to see collaboration and holistic thinking at every level – within each project team, between the six projects, and across local authorities – to maximise wider benefits, take advantage of a moment of change, and better inform development.

“ All of these are siloed projects, right? Machines learn from one another. So if you're creating a wonderful technology, but your underlying shared foundations for all of these are not working, they are only ever going to touch the surface of what they can do. So look holistically at the tech. - Participant

What role should the University of Cambridge play?

In discussions around the University's involvement in AI in local government, participants said they particularly value the independence of the University from both government and from private companies with financial motives.

Participants appreciate the University's collaborative approach, which brings together researchers, local authorities and the general public, to help ensure AI in local government is being developed effectively, iteratively and ethically.

They also welcome how this collaboration unlocks project funding capable of directly benefiting local residents. At the same time, participants point out that stakeholders outside the university provide a necessary counterbalance to academic priorities, and bring essential breadth and 'real-world' perspectives to R&D. Participants value the depth of expertise the University brings and feel this places the region and their local authorities at the forefront of advancements in AI.

SECTION 7

The people's ethical design framework for AI in local government

During the first phase of the public dialogue, participants discussed in depth their expectations of local government in the design and deployment of AI. In the second phase, one of the six live projects participants explored was an ethical framework for housing, structured around: public value creation; ethical and responsible innovation; governance and management; and data and digital infrastructure. Inspired by this framework structure, during the second in-person workshop, participants developed their own ethical design framework for AI in local government more generally. This section brings together participant recommendations from both phase one and phase two.

PUBLIC VALUE CREATION

Measurable and tangible benefit to the public

Evidence of the benefits brought by AI is essential to building trust. Participants call for clear measurements to monitor the contribution AI makes, with data on council services before and after AI's introduction. Importantly, they stress the need for qualitative data on how the introduction of AI has delivered real improvements to people's lives. In addition to data, over time participants would like to feel tangible benefits, for example with cleaner streets, better public transport or shorter waiting times for services.

“ A good starting point to get people on board and get people's trust is looking at things that are going to impact people, like public transport or getting you to school, not things that people don't think about. - Participant

AI addressing the right problems, where it can make a difference

AI in local government should directly address public needs and frustrations, as well as council priorities. Participants understand that local government is in the early, experimental stages of developing AI tools. But AI must provide a route to improving delivery of services, not simply provide more data in areas where the resources do not exist to action improvements.

Reinvestment of savings to improve services

Participants anticipate significant efficiency gains as a result of introducing AI, but efficiency in itself should not be the end goal. Time and money saved should be reinvested back into staff and services, to free up humans to do what machines cannot, from empathetic decision-making to understanding complex local contexts. As well as improving services, reductions in repetitive and mundane tasks should improve staff wellbeing and working conditions, not add to workloads or create pressure to achieve more.

“ If they do find savings, what they need to do is move those resources to areas of weakness within the Council. - Participant

“ Yeah, so say social housing, right, you may get an AI process to improve things and they said, OK, we've saved two manpower equivalents, shouldn't those two manpower equivalents be transferred to frontline services? - Participant

Inclusion and accessibility: public value for everyone

Participants feel strongly about the importance of inclusion in the broadest possible sense. They would like to see all council AI tools developed to include neurodiverse, vulnerable or disabled people, providing as wide a range as possible of formats and options for interacting with councils. Some also see great potential for translation tools to be integrated into AI tools. Participants would like to see measures in place to counteract digital exclusion, including clear communication about how public-facing AI tools work, so people know they exist and know how to use them.

Several spoke about the importance of carefully designing AI systems and human oversight procedures that account for unusual or individual situations, rather than treating people as part of a majority or narrow category. Most importantly, many participants emphasise that the simplest way to ensure inclusion and accessibility is to enable easy access to a person. This option should be built into any public-facing AI systems, and AI tools should never act as a barrier to human contact.

“ So you might have a very fantastic application. But if not many people are using it then it's not really serving the purpose... how can you reach public effectively? It's not just developing something useful, it's useful only when everyone can use it. - Participant

Social interaction and community before (or alongside) technology

Many participants do not want to see AI systems further reducing opportunities for human interaction, particularly for younger, older or vulnerable people who may already be suffering from social isolation. Participants emphasised the importance of the public feeling valued and cared for, rather than kept at arm's length. They discussed ways that technology has already undermined social life and community in the name of cost cutting or efficiency, such as automation in call centres or digital reception desks. Some participants appreciate the efficiency, independence and anonymity of carrying out basic tasks using screens, but many would like to see AI tools designed to bolster and enable human interaction, not replace it.

AI improving systemic issues, not replicating or papering over problems

Discussions often revolved around how AI tools could address underlying or broader systemic issues with council services. Participants would like local authorities to seize the opportunity of bringing in new AI tools to rethink and streamline, not replicate inefficient or outdated systems. A recurring theme amongst participants is the fragmentation of council services, which means residents being passed from pillar to post, vulnerable residents having to repeat procedures, information getting lost and services being delayed or failing to meet needs. Some can see great potential for AI to help connect services so they can 'speak to' each other, as long as it is not papering over inadequate foundations. In this potential state of flux and improvement brought about by AI, they would also like those designing AI tools to always think about how their tool could contribute to addressing wider systemic issues that matter to them, such as the housing trajectory connecting better with infrastructure and public service provision.

“ Is it going to improve the current system or just parrot the current system? - Participant

“ With AI, can it simplify anything so they don't have to fill in all the forms and they don't need as many people to answer the phones for instance. Could it help those kind of processes? So it's easier for the public to use council services or contact.

Because currently you go from one department to another. Is there a way that they could make it a central? Where they have all the information rather than all these different parts. - Participant

ETHICS AND RESPONSIBLE INNOVATION

Ask whether AI is necessary for the task

Many participants have grown much more aware of the potential trade-offs of introducing AI, and they would like any new application to weigh its potential social, economic and environmental costs. Several participants said councils should look at what already exists in the way of IT systems or human roles, to consider whether the best way forward could be to improve or adjust those, rather than being swept along by AI hype and applying it where existing solutions may be more suitable.

Protect and create jobs

Participants feel very strongly that from an ethical point of view, concrete measures should be put into place by local authorities to protect and create jobs, for those currently working at the council, and particularly for young people looking for work. Many participants do not accept job losses or resource constraints as inevitable, and feel that the high workloads and backlogs could be addressed with new roles as well as or instead of with AI. Many participants see the potential for AI systems to lead to job creation in roles such as editing, monitoring, auditing, cyber security and inclusion.

“ There's so much scope for positions to be opened up to a person, not a computer. Everyone talks about job losses, but they don't talk about the fact that there could actually be a job created instead. (...) like on say that the map team, there could have been a job created for someone else because they've got that backlog and they can't quite manage your workload as well. - Participant

“ If you've been working somewhere for 25 years and then they say AI has come to come to do your job and you're like, well, what, what do I do now? Do they have a path for sort of people? I think there's got to be a bit of a duty of care if you're going to be replacing people's jobs with AI. - Participant

Retraining instead of deskilling

Participants feel strongly about the importance of human oversight in roles where AI is used in public services. Time for human judgement within their role should be valued, properly resourced ring-fenced, to prevent over-reliance on AI, automation bias, or loss of skills and capabilities. Council staff should receive quality ongoing training to transition to or perform any new AI related roles well. 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.

Actively identify and remove bias

One area where participants see potential for job creation

is in tackling biases that AI could amplify or entrench in local government. AI tools should be developed and tested by diverse teams, and staff should be trained to spot and rectify bias as part of ongoing improvements. Some participants are particularly concerned about bias when historical data concerning vulnerable people is used to develop AI tools, such as in the bidding behaviour project. One participant suggested employing those with lived experience as well as experts such as sociologists, to advise on countering bias which could continue to disadvantage certain populations.

Ensure choice and consent

The wide variety of attitudes and approaches to AI amongst participants points to the importance of choice and consent over whether and how they interact with AI to access essential services. Many participants emphasise the importance of local authorities informing the public about exactly where and how AI is used and providing people with as much choice as possible. Several participants pointed out that some people may not want to interact with AI, others are hesitant about sharing their data with AI systems, whereas others see great potential in AI using their data in new ways to personalise services to them.

“ We're forcing people to engage technology whether they like it or not, people might not want to or it might affect their mental health or something (...) could there be an option just to not have it, to say I don't want any part of it? - Participant

GOVERNANCE AND MANAGEMENT

A Charter for AI in local government

Participants call for clarity around AI policy in local government. Some suggested this take the form of a binding charter – built on consistent key principles that could not be altered by an incoming political party. Participants would like guidelines to be in place to ensure services are high quality, ethical and secure, consider the broader impacts of AI, and prevent mission creep or profiteering.

Independent regulation strengthened by involvement of all stakeholders, especially a broad public

Some participants would like to see an independent auditing and regulation body, with the power to enforce standards. But they also expressed their distrust of existing independent regulators such as Ofwat or Ofgem, which they feel have been ineffective and unaccountable to the public. Many participants suggest strengthening independent regulation with a democratic approach, involving committees of stakeholders including representatives from across council departments, researchers, AI developers, and importantly members of the public with a broad range of backgrounds, ages and needs.

“ Is there a need for an independent body to actually regulate and audit local government use of AI? So it's not the local government marking their own homework. That actually has some regulatory might, like... it's got some teeth. - Participant

“ And like a committee. With all the areas of the council, and unified, democratic decisions, not one person saying what we should be doing... individuals from local wards, not aligned to a political party, people with lived experience who are paid so there's no disadvantage... an ongoing committee, and people have a voice in how it develops. - Participant

End to end transparency

Participants feel the public needs to know from the outset how particular AI applications are chosen to be used in local government, who is involved in developing, supplying and hosting them, and what data is being used to train them. They would like to know exactly where and how AI is being used in public services, how their data is being used and stored, and call for decisions or conclusions made involving AI to be traceable and explainable, ideally by staff involved in related public facing roles. Participants agree that information about value for money and the cost of AI compared to the cost of staff completing particular tasks should be made easily available, including information about where and why savings from the use of AI have been reinvested. To build trust, participants would also like to hear about when things have gone wrong or not worked out as expected, and what has been done about it.

“ There should always be a place on the website or somewhere that shows you which part of a process is AI. Because if something goes wrong, you can understand that's potentially human error or is that a computer error? - Participant

Human accountability and clear routes to redress

Participants agree that a named human must be responsible for any decisions made or actions taken involving the use of AI in local government. There need to be clear lines of accountability and routes to redress if residents feel mistakes have been made.

“ You need a human to make the final decision, you can't hold AI responsible can you. You can charge it with any mistake or any wrongdoing or any fraud or whatever. - Participant

Clear, accessible reporting, actively communicated to the public

Participants say governance information should be communicated clearly to the public and made easily available, rather than being buried in meeting minutes or lengthy reports

“ There needs to be transparency on reporting issues as well. Like I know from experience, it's very easy to bury things in a council agenda because no one ever reads it and we have no local journalism anymore. So, you know, some of these things only get found out because people are actually going to read 100 page reports. - Participant

“ Like you'd have a corporate responsibility report. I feel like there should be an AI report. - Participant

Future proofing: 'hope for the best but plan for the worst'

Participants say 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.

Public engagement throughout the process

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.

“ Reach out to working people, reach out to unemployed people, able people, disabled people. You go across the board, single parents, nuclear families, all of them. Yeah, you find out the needs and you create, you know, from that you streamline from what are the key concerns. - Participant

“ More public dialogues like this. - Participant

Communication about AI in local government: meet the public where they are

As well as public engagement, participants call for active communication about all aspects of AI in local government, in traditional and new ways. As well as clear information available on council websites, participants suggest public outreach at local events and including information on council tax bills about time or cost savings that have been made and how they have been reinvested to improve services.

DATA, DIGITAL AND INFRASTRUCTURE

Ensure an ethical supply chain and prioritise local developers

Participants expect AI researchers and developers to work directly with local authorities, rather than councils buying in systems developed by unaccountable tech giants. Information about procurement choices and who is involved at every stage of the design and deployment of new AI systems should be in the public domain.

Design, develop, train and test with diverse teams

To respond to public needs, and to counteract existing systemic bias and the potential for AI to amplify those biases, participants expect new AI tools to be developed and tested by diverse teams of people, including those with a range of lived experiences and expertise in inclusion and anti-bias. Political bias is also important to act upon as participants agree AI should not reflect the political allegiances of any local council.

Store data in the UK

There is clear agreement that participants expect data to be stored and processed locally within the UK, and kept under UK legal jurisdiction, as opposed to using cloud - based storage or opening up the possibility of data sharing outside the country.

" I am concerned about the role of data, look at the NHS and what we're doing with that American company. I do not want that happening with my data in local government. And I think that needs to be written in early on in the into the guidelines because at some point in the future an American company will come in and they will be better and cheaper. - Participant

Plan, test and monitor robustly, for agile, ongoing improvement

Participants emphasise the importance of long-term planning and stringent testing before any new AI systems are rolled out. Participants feel AI should be introduced incrementally in local government services. They want to see systems proven to work well in lower-risk applications, such as the MAPLE, SHIFT or Fly Tipping projects, 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. They emphasise the need for a flexible and agile approach to make constant improvements in this new, fast-paced and ever-changing field.

Measure and minimise the environmental impact

Participants are increasingly aware that data centres and other infrastructure necessary for the AI they are familiar with, such as ChatGPT, are highly energy intensive and require large volumes of water to function. In the context of local authorities Net Zero goals, participants call for transparency over the climate impacts of AI, as well as measures to mitigate them. For example, by being selective over where AI is used, how and

how much data is stored, and prioritising AI projects with smaller energy and water footprints, such as task specific applications.

Accuracy, traceability and correction

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.

Data security and system backups should be well resourced

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. A risk-aware approach should be fundamental to design, not an afterthought and security measures and backups against both malicious attacks and system failures should be in place. These security measures extend to the staff handling sensitive personal data, who should be appropriately vetted. Participants also say the public should be protected from the increased possibility of scams and cybercrime if AI becomes part of their interactions with their council.

Privacy: data use and access must be clear and boundaried

Participants expect the uses of their data to be made clear, and their data should only be made available to those within the local council who need access to it to provide services. Several also expect to be able to access the information the local authority holds about them, and to be able to adjust, update or retract the information used as part of AI systems.

" I want to have autonomy over my own information. We need to have trust that our government have the capabilities to use it in the right way, and however you're using my data is being done in an efficient, secure manner, and that they're not going to sell it. - Participant

Use every opportunity to improve, join up and standardise local authority IT architecture

Participants see great potential for AI to connect fragmented departments and systems. AI should not simply be working hard to compensate for shaky foundations, but its development and deployment should inspire and enable change to bring about more efficient, consistent and centralised structures.

SECTION 8

Conclusion

The public dialogue described in this report has been an ongoing conversation about AI in local government between Cambridgeshire residents, researchers from the University of Cambridge, and local government representatives.



It began with residents exploring in depth their hopes, concerns and principles for AI in local government. Their insights fed into six collaborative projects, which paired University researchers with local councils, to develop and deploy practical AI solutions for public services in line with residents' views. These projects were then presented back to residents in a second public dialogue phase, where participants discussed with project leads their hopes, concerns and recommendations, before bringing together their deliberations to develop an ethical framework for AI in local government. Looking ahead, the AI for Local Government Accelerator will continue to support the six funded projects as they develop and evaluate their approaches in partnership with local authorities. The findings from this dialogue and the ethical framework developed by participants will help shape the next stage of project development. Lessons from the programme will also inform future collaborations between researchers and local government.

Participants would like to see AI deployed selectively and thoughtfully in local government, in ways that tangibly serve their needs, as well as the needs of local councils. They see great potential for AI to make public services more efficient and joined up, but believe every deployment decision should be weighed against its wider societal impact.

AI should not replace humans or be used to justify cost-cutting, but enable staff to serve the public better. Public facing AI tools should be user friendly, accessible and inclusive, actively tackling bias.

Trust must be earned by demonstrating the public value AI delivers. Clear guidelines and transparent, independent governance should ensure ethical and local design, procurement and deployment. Robust ongoing monitoring should safeguard and improve accuracy, data privacy and security, while minimising environmental impact.

This dialogue makes clear that residents are willing partners in developing AI for local government, and that ongoing public involvement is essential to ensuring deployment meets the public where they are, meaningfully contributes to the issues they care about, and improves people's lives.



SECTION 9



Acknowledgements

ai@cam and HVM would like to thank all the participants in this Public Dialogue on AI in Local Government, whether they shared their views online, in-person or both. Their thoughtful and committed engagement with this topic has provided valuable new insights to inform the design and deployment of AI in a way that truly serves the public interest.

We would like to thank the Leads from the six projects selected to join the AI for Local Government Accelerator programme, for presenting and discussing their projects with clarity and openness; Alexis Litvine and Guillaume Proffit on MAPLE: Map Automation for Planning and Local Efficiency; Matteo Zallio and Anya Petyaeva on SHIFT: AI Enabled Surveys for Housing Intelligence & Future Trajectories; Tyler Holderness on Deep Learning for Fly-tipped Waste Detection; Jerry Chen on Bidding Behaviour in Choice-Based Lettings; Ronita Bardhan and Adhib Syed on PRISM: Predictive Risk Intelligence for Social Housing Maintenance, and Viviana Bastidas Melo on Human Oriented AI: Design framework for reaching Vulnerable Tenants.

Thanks to our workshop one speaker Jennifer Schooling from Anglia Ruskin University; to Tom Bradshaw from Scriberia for providing us with a visual record of workshop two, and to all observers and specialists who joined us in the in-person workshops from local government and from the University of Cambridge.

Thank you to our report authors from Hopkins Van Mil: Henrietta Hopkins, Director; Pauline Harris, Deputy Director; and Emma Rigard-Asquith, Researcher.





**For questions about this
report, get in touch via
the information below:**

Email:

contact@ai.cam.ac.uk

Visit:

ai.cam.ac.uk

Connect on X:

@ai_cam_mission

Connect on Bluesky:

@ai.cam.ac.uk

LinkedIn:

@ai-cambridge

In partnership with:



hopkinsvanmil.co.uk