

Tiina Härkönen, Kirsi Hantula and Marko Aalto

PARTICIPATORY AI DEVELOPMENT MODEL

A citizen dialogue on the use of AI
in public services

Sitra memorandum

© Sitra 2026

Participatory AI development model

A citizen dialogue on the use of AI in public services

Authors:

Tiina Härkönen and Kirsi Hantula from Sitra, and Marko Aalto from Reaktor.
Tiina Härkönen and Kirsi Hantula are Senior Leads in Sitra's Democracy Innovations programme. Marko Aalto is Vice President of Data and AI at the technology consultancy Reaktor.

Working Group:

Sitra: Tiina Härkönen, Kirsi Hantula, Tarmo Toikkanen, Antti Lehtinen. Reaktor: Marko Aalto, Anne Karumo, Akseli Kouvo. VTT: Anton Sigfrids. Digital and Population Data Services Agency: Marko Latvanen. Veikko Isotalo (University of Turku)

Acknowledgements: Reijo Aarnio, Sitra.

Layout: Grano Oy

ISBN: 978-952-347-461-1 (PDF) www.sitra.fi/en/

ISSN: 2737-1034 www.sitra.fi/en/

Sitra memoranda are insights produced to support our future-oriented work.

Contents

Foreword	4
Summary	5
1. Introduction: The lack of trust in AI calls for action	7
2. Case example: What do you think about AI, Finland?	
– A national dialogue	11
3. Launching a digital citizen dialogue and the invitation process	13
4. Delimiting the Voxit dataset for the creation of citizen rules	22
5. Preparing the citizen rulebook	26
6. Citizen rules for AI in the development and use of public services (Finnish experiment)	29
7. Conclusions and closing remarks	35
Appendices	38

Foreword

Artificial intelligence is becoming an increasingly influential part of society, with its impact extending into everyday life, work, services and decision-making. Sitra's *Megatrends 2026* report highlights how AI is challenging, among other things, traditional approaches to education and decision-making. While development offers many opportunities, it also raises questions about how AI can be used responsibly and purposefully. Whatever one's views on AI, the changes it brings will play a significant role in shaping the years ahead.

In this study, we provide one important viewpoint for the broader discussion on AI: insights from the 6,500 people who took part in the national dialogue *What do you think about AI, Finland?* The aim of the dialogue was to deepen understanding of what kinds of thoughts, expectations and concerns people have about the use of AI in public services. We did not seek to produce definitive answers or recommendations, but rather to identify the key questions and themes shaping public perceptions of AI in everyday life. We also wanted to offer a widely representative overview of the value-based public debate across different parts of Finnish society.

AI is developing so quickly, and its effects are often difficult to grasp. Many people have strong opinions and feelings about AI, even though these may not be based on personal experience. Public sector decision-makers need a clear understanding of how people perceive the opportunities and risks related to AI to facilitate responsible and trustworthy governance.

The insights highlighted in this report show what themes are important in the Finnish AI discussion, what people hope for, what worries them, and what expectations they have – expressed in their own words. The study brings these diverse viewpoints together and identifies shared themes that can serve as a foundation for ongoing public debate and future responses.

We hope that the findings will support the national conversation on the role and future of AI. Discussions about AI require voices from the entire society. We would like to thank the technology consultancy Reaktor for their valuable contribution to the preparation of this report, as well as the citizens and organisations who participated in the national dialogue. Without their involvement, this work would not have been possible.

Kalle Nieminen

Director, Solutions

Sitra

Summary

Artificial intelligence has rapidly become a part of everyday life in Finnish society, yet public trust in those who develop and use it remains low. In the development of generative AI, individuals' sensitive data forms a significant resource, meaning that citizens should have genuine opportunities to influence the direction in which AI is developed and how it is used. This is particularly true in the public sector, where decisions have broad impacts on daily life.

Trust in AI is built on transparency, inclusion, and the recognition of citizens' epistemic rights. The Finnish Constitution states that citizens must have the possibility to influence decisions that concern them. This right should also extend to the development and deployment of AI. Finns would like to influence decision making and shared societal matters easily and quickly online, but the discussion culture on social media discourages many from participating.

Without transparency and genuine agency, sustainable trust in AI cannot be established. Public administration and other societal institutions should view the safeguarding of citizens' AI agency as a core responsibility. The framework of epistemic rights summarises citizens' rights in the age of AI into four areas: access to information, access to skills and capabilities, participation, and the right to form opinions free from pressure. Strengthening these rights is essential for maintaining trust and legitimacy.

The model for participatory AI development and the technical citizens' AI rulebook – developed through bold experimentation – are examples of how citizens' perspectives can be incorporated into AI development and decision-making.

The citizen dialogue *What do you think about AI, Finland?* organised by Sitra and Reaktor, was carried out in September 2025 on the anonymous, constructive discussion platform Voxit and engaged nearly 6,500 participants. The discussion published 194 statements about AI development and use, of which 169 were created by citizens.

The aim of the dialogue was to address challenges of trust from the citizens' perspective and to create a rulebook for the use of AI in public services. The discussion sought to increase public understanding of AI, give citizens agency, and test and refine the rules created by citizens. The goal was to produce a technical rulebook for guiding a language model and to document the entire operating model in a way that can be utilised elsewhere in Europe.

The dialogue produced a rich and diverse dataset of citizens' views on the use of AI in public services. Based on this material, a technical rulebook

was created to influence the operation of AI, and it is presented in this memorandum. The result is a valuable foundation that supports the responsible and democratic use of AI in society and helps to identify solutions for safeguarding public trust.

1. Introduction:

The lack of trust in AI calls for action

Artificial intelligence is becoming increasingly embedded in everyday life yet trust in its developers and use remains low (The University of Melbourne, 2025. *The Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025*). Traditionally, trust has been viewed as essential for the adoption of new technologies.

It can also be argued that because data generated through ordinary people's online behaviour has been turned into raw material for generative AI, individuals should therefore have a genuine opportunity to influence the direction of that development. Trust is built through transparency, participation and the recognition of citizens' epistemic rights.

We are living through a transitional phase in which society has already, in part, entered the era of AI. Public debate on the topic is mainly dominated by experts. The moral and societal choices linked to AI development, meanwhile, seem to be made almost exclusively by a small group of global technology giants as they produce AI models for worldwide markets – including the public sector. The voices of ordinary people have been left virtually unheard.

Our behaviour, networks and other sensitive personal information are used, among other things, to train generative AI models. When the data invisibly collected from individuals forms a critical foundation for AI development, it is only logical that people should have the possibility to influence how that data is used and in which direction the technology evolves. While this may be an unrealistic expectation for the largest AI developers, the situation could be different for public sector actors. AI affects us in two distinct ways: beyond training, it is increasingly used within digital systems to process information about us, alongside, and as part of, traditional IT.

In Finland, the Constitution states that citizens must have the opportunity to influence decisions that concern them. This right should also extend to the choices made during the development and deployment of AI, particularly when these decisions significantly shape public services and, through them, everyday life.

Although studies (Sitra, 2022. *Demokraattiset osallistumismahdollisuudet Suomessa*) show that citizens' trust in generative AI and its developers is low both internationally and in Finland, new technologies still have the potential to improve public sector productivity. Due to

demographic changes, labour is becoming scarcer relative to need, meaning that technology also has the potential to compensate for the emerging resource gap.

AI research further shows that trustworthy AI matters: when people consider AI systems reliable, they are more likely to accept their use.

Studies (Sitra, 2022. *Demokraattiset osallistumismahdollisuudet Suomessa*) also indicate that Finns would like to influence decision making and societal matters easily and quickly online. However, widely accessible social media platforms tend to amplify the loudest and most extreme voices, discouraging many from participating. To build social trust, people must be actively involved in the development and direction of services, both as citizens and as customers and users of services. Without transparency and genuine agency, people will not trust AI in the future either.

The Nordic countries have a long tradition of liberal democracy, and societal trust is generally high. Defending and developing democracy, however, requires continuous, open and active dialogue with citizens, and this dialogue must also extend to debates on AI. Listening to citizens becomes especially important when we discuss publicly funded services. Automated decision making is already used in various forms, and the increased integration of AI into services is a natural continuation of digitalisation.

Although not everyone is an AI expert, each person is an expert in their own everyday life and can assess how they wish to be treated within public services.

Epistemic rights as a foundation for strengthening citizens' agency

For citizens to genuinely influence the choices made in public sector AI development, public administration and other key societal institutions must increasingly recognise the safeguarding of citizens' AI agency as part of their essential role. They must understand that protecting citizens' AI agency is a precondition for ensuring that public services, which are increasingly powered by AI, remain trusted and legitimate in the eyes of the public.

In the rapidly evolving AI era, safeguarding citizens' agency requires more purposeful action and more strategic coordination from public sector actors and others. All organisations need to recognise how they can strengthen citizens' AI agency naturally within their own remit, using the resources available to them.

The framework of epistemic rights, published by Sitra in 2025, provides support for this work, summarising citizens' epistemic rights in the age of AI into four essential rights (Sitra, 2025. *Tiedolliset oikeudet tekoälyn*

aikakaudella). The original framework, adapted in this memorandum for the context of AI development, was drafted by Emeritus Professor of Communications Hannu Nieminen.

This memorandum describes concrete ways in which public administration, the media, and the education sector, for example, can strengthen these rights.

By investing in the protection and reinforcement of these four epistemic rights, a wide range of organisations can meaningfully work towards a common goal. The framework provides a shared, value based foundation built on the understanding that in societies where the role of AI is rapidly expanding, people’s opportunities to influence and exercise agency must be protected more vigorously. At the same time, the rights offer a shared point of reference through which different actors can develop practices and operating models that strengthen citizens’ AI agency.

Figure 1: Epistemic rights strengthen citizens’ agency in the age of AI

ACCESS TO INFORMATION

All individuals must have equal access to essential and accurate information about the development and use of AI.

SKILLS AND CAPABILITIES

All individuals must be guaranteed equal access to the capabilities needed to understand the opportunities and challenges of AI development, enabling them to protect their personal or societal interests based on this understanding.

PARTICIPATION

All individuals must be granted equal opportunities to influence the value base, direction and impacts of AI development through public discussion and decision making.

NO COERCION OR PRESSURE

All individuals must have equal opportunities to form views about emerging AI without coercion or manipulation by actors who hold significant epistemic power.

The framework of epistemic rights summarises the measures that strengthen citizens’ epistemic agency into four core rights.

The concept of epistemic rights will not become established in society unless it is supported by dedicated advocates who systematically work to strengthen them. Public sector organisations, which are increasingly adopting AI, are naturally part of this group of advocates.

These organisations should review their own operations to ensure they promote all four epistemic rights. Public administration is obliged to produce and provide accessible materials about AI and support citizens' AI literacy. By creating and adopting new operating models that strengthen citizens' agency, public sector organisations can ensure that the boundaries of AI use in public services are subject to public scrutiny. Public administration can also support the dissemination of epistemic rights by ensuring that its own AI guidelines, recommendations, programmes and policies comply with the principles of epistemic rights.

In line with the third epistemic right, all individuals must have equal opportunities to influence the value base, direction and impacts of AI development through public debate and decision making. The participatory AI development model presented in this memorandum provides a framework through which public sector actors can put this principle into practice. The technical rulebook is one means of translating citizens' views into concrete form so that they can be genuinely considered in the development of AI systems and related decision-making in public administration.

The purpose of this discussion and memorandum is not to provide legal analysis or propose new legislation governing the use of AI, but to present citizens' views, understand the challenges surrounding the topic, and encourage debate on solutions that help build trust.

2. Case example:

What do you think about AI, Finland? – A national dialogue

The adoption of AI is not a black-and-white issue. It is not simply a question of whether individuals and organisations adopt AI or not, but rather a collection of different attitudes and approaches that form a *continuum of adoption* – from complete non-use to advanced utilisation.

Some studies suggest that much of the current use of generative AI is quite superficial. Generative AI is used, for example, for drafting emails or summarising reports. Although these uses bring only marginal benefits at first, they familiarise people with the tool and enable deeper and more advanced applications later.

If the goal is to increase the number of people using AI, trust is needed alongside skills. The aim of the national dialogue *What do you think about AI, Finland?* was to address exactly this challenge from a citizen perspective by creating a trusted rulebook for the use of AI in public services.

At the same time, decision makers, civil servants and AI service developers sought insights into potential problem areas and a clearer picture of citizens' readiness to accept and engage with AI.

Objectives of the dialogue

- 1. Increase public understanding of AI** by exploring challenges, risks and opportunities from a citizen perspective.
- 2. Strengthen citizens' agency** by establishing a nationwide, cross societal digital dialogue on AI.
- 3. Increase trust** in the development and use of AI in public services by creating citizen-generated rules for the use of AI in the public sector. These rules cover the design, governance and technical implementation of public sector AI systems.
- 4. Test technical rules** applied to a language model to influence its functioning.



5. **Create a prototype library of system prompts**, enabling language models to be steered in ways that reflect citizens' values and expectations.
6. **Document a citizen engagement model** that can be replicated and used in any European country.
7. **Deliver the results to decision makers and civil servants.**

The Finnish experiment was designed for a broad audience in the form of a public discussion on rules for AI use. This inevitably means that producing precise instructions and rules based solely on discussion data is challenging. It must therefore be recognised that interpretive authority shifts elsewhere, but also that many citizen-generated statements are valuable in ways that extend beyond being mere rules for a language model.

This memorandum describes the AI discussion experiment, the participatory AI development model built from it, each stage of the process, and the resulting citizens' AI rules. Each stage includes practical guidance on what successful implementation requires, as well as suggestions on how the model could be further improved.

THE PARTICIPATORY AI DEVELOPMENT MODEL IN A NUTSHELL

- The participatory AI development model is a method designed to help European democracies strengthen human-centric and citizen-driven AI development in public administration.
- In the model, citizens participate in a nationwide digital dialogue about AI. In the Finnish pilot, the dialogue took place on the Voxit digital participation platform.
- Implementing the process requires a minimum duration of 3–4 months.
- Sufficient resourcing is essential, for example, to support communication around the citizen dialogue and to ensure high quality moderation. The team implementing the model should include expertise in digital participation, technical understanding, and ethical AI development.
- It is advisable to ensure early in the process that a public sector actor is committed to using the resulting AI rulebook in its service development. In Finland, such a partner was not yet available due to the experimental nature of the project.
- To build trust, transparency is crucial: both the interpretation and analysis of results, as well as the Voxit result report, must be openly accessible.

3. Launching a digital citizen dialogue and the invitation process

A digital citizen dialogue is appropriate when insights from large groups of people are needed, and when, for example, decision makers and civil servants require an up-to-date picture of an issue that is important from the citizens' perspective. In such cases, there is often also a need to build broader mutual understanding between citizens and representatives of public administration.

The use of generative AI in public services was a suitable topic for a citizen dialogue because it affects almost everyone's daily life. It is also seen as a key technology for the development of services. AI must, however, be used in ways that serve citizens and remain fair, so that trust in public administration is not undermined.

It is essential to understand where mistrust toward AI and its developers originates, and how it is directed, in order to identify and address the underlying challenges.

These needs provided the basis for the *What do you think about AI, Finland?* experimental project, in which Finns were invited to engage digitally in discussion on the risks and opportunities of AI in public sector services.

A promise to participants

The citizen dialogue was created to deepen understanding of the needs, expectations and concerns Finns have regarding AI – issues they wanted public sector decision makers to understand.

Participants were told that the dialogue would generate material for creating a technical rulebook for AI. The material would be used as training input for a language model, allowing the rulebook concept to be tested. No comparable AI rulebook based on ordinary citizens' views has been produced anywhere in the world.

Participants were also informed that the results would be brought to the attention of decision makers and civil servants, and that the longer-term goal could be to develop an AI model for public services that operates in line with citizens' perspectives.

Choosing the method

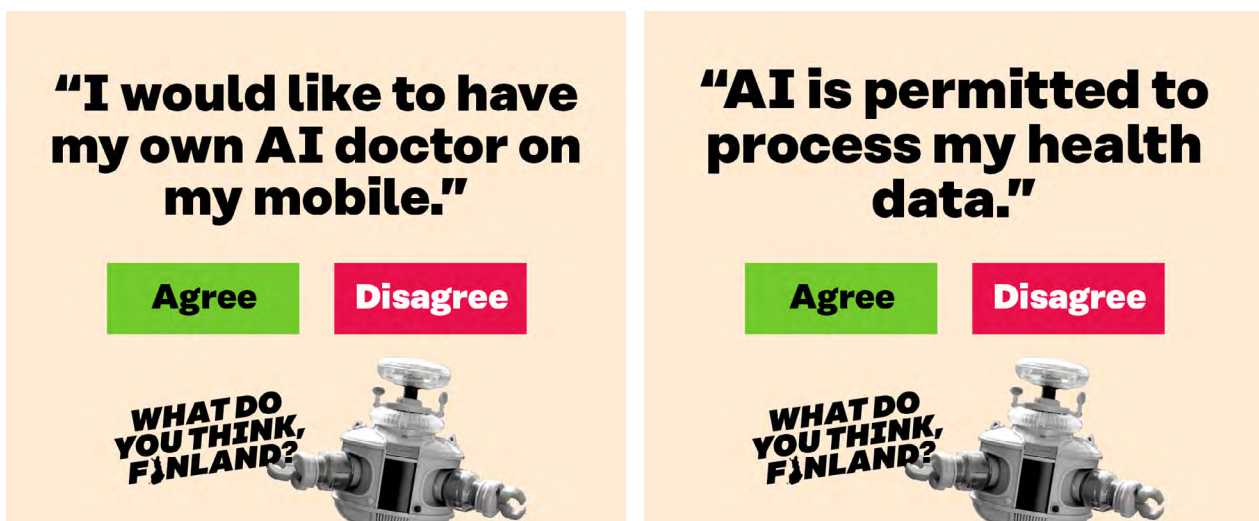
From the outset, the design of the AI dialogue aimed to create a model that could be replicated elsewhere whenever citizens' views need to be heard in line with democratic principles.

When broad participation is desired, barriers to taking part must be minimised in every way. Speed, ease of participation and independence from time and place are crucial. This naturally points to digital democracy technologies designed for constructive public debate, such as digital participation platforms.

Participants must also be able to trust that they can join the societal discussion safely and without interference, making anonymous participation essential. Online harassment can be prevented by choosing a platform with built-in moderation features.

For the Finnish national dialogue, the platform chosen was Voxit, a digital participation tool designed for large-scale citizen engagement. Due to its operating principles and moderation features, it does not allow the types of targeting, harassment, participant identification or silencing of moderate voices that are common on social media. The discussion remains genuinely constructive and appears solution oriented and respectful to participants.

Figure 2. Ads from the *What do you think about AI, Finland?* citizen discussion.



Finns were invited to share what risks or opportunities they see in the use of AI in the digital public services used by citizens. People were directed to present their views on Voxit, a digital participation platform designed for constructive public discussion.

WHAT IS VOXIT?

Voxit originated in the United States in the mid 2010s, when the civic organisation *Computational Democracy Project* developed an online platform called *Polis* for constructive citizen deliberation. Voxit is an open-source solution that has been developed in Finland as a dedicated fork of the original Polis platform. It has been adapted to meet the requirements of the EU's General Data Protection Regulation (GDPR) and European accessibility standards. In Finland, particular attention has been given to mobile usability. The name Voxit comes from the Latin word *vox*, meaning 'voice'.

The Voxit platform provides real-time data on the distribution of participants' opinions. Results show which statements the majority agree on and where opinions diverge. Using statistical analysis, the platform also clusters participants into opinion groups based on their voting behaviour. In this way, the results improve both citizens' and decision makers' understanding of the issue and strengthen the knowledge base for decision-making.

A key feature of the platform is that participants can vote on statements, while at the same time proposing their own statements for others to vote on, ensuring that perspectives they consider important are represented. A typical Voxit discussion is open for about a month, and participants can propose new statements during the first half of the discussion. In the latter half, only voting is possible.

Inviting citizens to the Voxit dialogue

In the Finnish experiment, people were reached using several methods: paper invitations sent to a random sample, email newsletters and stakeholder communications, social media posts, and digital advertisements on tabloid news websites and social media. The dialogue was open to everyone, and Finnish citizenship was not required. Participants could participate in the discussion in Finnish, Swedish or English. Voxit includes a translation tool, allowing statements to be added in multiple languages and viewed in the participant's own language.

The invitation process did not aim for perfect statistical representativeness of the Finnish population; instead, it sought to recruit as many participants as possible.

INVITATION PROCESS FOR THE *WHAT DO YOU THINK ABOUT AI, FINLAND?* DIALOGUE

1. Representative sample and physical letter

- A random sample of 10,000 residents, including postal addresses, was purchased from the Finnish Digital and Population Data Services Agency (DVV). A paper invitation was sent before any other communications, containing a QR code for easy mobile access. The invitation was in Finnish and Swedish.
- The letter also provided instructions on how to use Voxit with assistance, ensuring that even individuals unfamiliar with digital devices could participate. Because participation is anonymous and does not require login, this can be easily managed.

2. Display and social media advertising

- Participants were targeted via display ads on major tabloid websites and other digital services within the same media group. Ads included selected eye-catching Voxit statements.
- Social media advertising used the same content and targeted different demographics across Instagram, Jodel, LinkedIn and Facebook. Paid campaigns were conducted primarily on Meta platforms and Jodel.

3. Other communication

- Decision makers and experts in business, the public sector and NGOs were reached through direct emails and posts on the social media channels of the organisations conducting the experiment. Event marketing and a panel discussion were also used for the launch.
- A press release announcing the launch was distributed through a media distribution service in Finland.

Statements on the Voxit platform

A Voxit dialogue begins by publishing 15–25 initial statements: background statements and seed statements. In the Finnish experiment, the platform was open for approximately 3.5 weeks.

Background statements

Voxit is not a survey tool and does not provide detailed sociodemographic data. Background (meta) statements allow organisers to capture basic participant characteristics and monitor who is joining. They can also be used for correlation analysis.

Example: *“I am under 20 years old.” / “I have a higher education degree.”*

Background statements take up space and provide little added value to participants, so their number should be kept low. In the Finnish experiment, five background statements were used.

Seed statements

Seed statements initiate the thematic direction of the discussion. Their formulation requires time, as they shape both the tone and the level of abstraction.

Example: *“I suspect that AI places too much emphasis on the interests of minority groups in public services.”*

In the context of creating a citizens’ AI rulebook, seed statements must guide participants to generate statements that are sufficiently concrete while not referencing specific use cases that would make rule generalisation difficult. Striking the right level of abstraction required extensive internal discussion and iteration.

It is essential to clearly document the principles followed when drafting the seed statements to demonstrate that the team sought to ensure fairness and balanced perspectives.

How seed statements were formed in Finland

A set of existing documents on ethical, human-centred and democratic AI development guided the formulation. These included:

- The EU General Data Protection Regulation
- UNESCO recommendations on AI ethics
- Sitra's epistemic rights for the AI era
- Helsinki and Tampere city principles for ethical data and AI use

The team was also aware of a 2023 US experiment in which a smaller group developed 75 rules for training a language model.

The Finnish team first merged recommendations from the ethical documents into a set of 13 overarching principles (Appendix 1). These covered fairness, non-discrimination, privacy, safety, explainability, support for citizens' AI literacy and participation, human oversight, responsibility and ultimate decision-making authority.

The team then attempted to create concrete seed statements under each principle. They began by mapping the 75 US statements onto the principles, leaving only a few unmatched. They then revised, rewrote or replaced statements to suit the Finnish context and spent substantial time refining the wording.

The iterative process produced 21 seed statements in total (Appendix 2). Because the team lacked sufficient shared working time, most work was done individually. Time pressure towards the end meant that the final statements were not fully assessed against the original principles.

Future processes should allow ample time for reviewing seed statements and ensuring alignment with the principles guiding their creation. The workload can be reduced by using previously tested statements as a reference.

A final evaluation step could mirror the method used at the start of the Finnish process: mapping each seed statement under the principles to confirm broad coverage of key ethical and human-centric AI themes, and to ensure that no principles have been unintentionally compromised during iteration.

Citizen generated statements

A core purpose of the Voxit platform is to strengthen citizen agency. This appears clearly in the approach to statements: the seed and background statements merely start the dialogue, while most statements are generated by the participants themselves. This gives participants significant influence over the direction of the conversation.

Participants often bring entirely new perspectives that organisers had not considered. Their statements can refine or challenge the ideas embedded in the seed statements.

In Finland, this dynamic was strongly visible. By the time the dialogue closed, the platform contained 194 AI-related statements, of which 169 were citizen-generated. Before the launch, organisers had agreed to cap the number of published statements at 200 to avoid overwhelming participants.

The dialogue ran for just under a month in September 2025, with citizens able to propose new statements during the first two weeks.

CHECKLIST FOR ORGANISERS PREPARING A VOXIT DIALOGUE

- Allocate sufficient time to draft seed statements clearly and at the right level of abstraction.
- Write statements as unambiguously as possible so participants interpret them similarly. The style chosen for seed statements will shape how formal or informal participants expect their own statements to be.
- Centralise all citizen dialogue information on a single webpage, including method details and participation guidelines.
- In the Finnish experiment, anyone living in Finland could participate without identification or creating an account. Access from outside Finland was blocked to limit interference.

Moderation of the dialogue

Because most citizens lack deep knowledge of AI or public sector automation and privacy regulation, such expertise should not be expected in a citizen dialogue aimed at developing AI rules. People do not need technical expertise to express views on the values guiding AI development or the limits that should be placed on its use.

In Finland, differences in AI knowledge were visible, for example in the difficulty many had distinguishing between generative and nongenerative AI. A narrower topic definition might have helped some participants, but organisers must accept that most people will not read detailed explanations even when available.

Differences in AI literacy inevitably affect statements. Moderators should expect some statements to contain factual errors, and that many participants will want to discuss AI broadly at the societal level. This was clear in Finland, where participants expressed strong interest in themes such as AI education and rules for AI procurement.

Accepted examples:

- *“AI will increase inequality because not everyone knows how to use or understand it.”*
- *“A European language model is a better choice for us because EU governance allows us to influence decisions.”*

Accepted factual concern:

- *“Healthcare visits and appointment details must not be recorded or used as AI training data without permission.”*

Rejected example:

- *“I don’t want AI; I’m 82 years old.”*

Moderators hold significant influence and must prepare carefully and act responsibly. Their personal values must not influence decisions. Complex or contradictory statements enrich the dialogue and reveal citizens’ concerns.

To maintain quality in a large scale national dialogue, moderators should define key moderation policies in advance and publish clear moderation rules (Appendix 2).

In Finland, moderation was relatively strict: nearly 400 statements were rejected, mostly because similar statements already existed or because they were too general for rulebook development. Only a handful were inappropriate.

Moderators also agreed that no individual subtheme should account for more than one fifth of all accepted statements, ensuring that no single issue, such as privacy, would dominate the conversation.

PROJECT IMPLEMENTER’S CHECKLIST FOR MODERATING A VOXIT DISCUSSION

- Moderating a Voxit discussion with thousands of participants requires sufficient resourcing and personal preparation. There should be at least two to three moderators, and they must be able to dedicate a few hours each day to the task.
- When planning the moderation of a Voxit discussion, it is important to remember that the need for moderation is temporary. Participants’ ability to propose new statements for the Voxit discussion usually ends around the halfway point of the discussion period. After that, active moderation is no longer required. The need for moderation can also be limited at the beginning of the discussion by stating in the moderation rules that moderation will take place only at specific times – for example on weekdays and during office hours.

- The moderation needs of a large national Voxit discussion can be anticipated in advance once the scheduling of the communication supporting the discussion is known. In practice, whenever more prominent communication is issued about the discussion, a spike in participant numbers occurs on the Voxit platform. When many participants visit the platform, they propose more new statements, which the moderators should review as quickly as possible.
- Despite thorough preparation, moderating a large Voxit discussion can at times be challenging. Even though clear moderation rules have been established for the discussion, they may still leave room for interpretation. In difficult cases, the best approach is to make these interpretative decisions together in discussion with the other moderators.

4. Delimiting the Voxit dataset for the creation of citizen rules

Once the Voxit discussion has ended, it is time to begin analysing the dataset collected on the Voxit platform. This work is carried out by specialists. The aim is to formulate rules, based on citizens' views, that support the governance and development of AI systems used in public services.

Because the set of statements obtained from the Voxit platform is likely to cover a very wide range of perspectives related to the use of AI, it should be anticipated during the analysis phase that only a small proportion of the statements produced by citizens can be incorporated into the rule set as they are.

Principles guiding the selection of the base statements for the citizen rulebook


The selection of the base material for the citizen rulebook begins by establishing the principles on which the choice of statements can be carried out transparently and systematically. As the aim is to create AI rules that genuinely reflect the shared values and views of citizens, it is equally important to ensure that the statements forming the basis of the rulebook have received sufficiently broad support in the Voxit discussion.

On this basis, the Finnish working group established two principles to guide their analysis and to define the boundaries of the dataset.

1. The Voxit platform generated two distinct opinion groups among the Finnish participants in the public discussion: group A, 868 people (13.5% of all participants), and group B, 3,799 people (59.1% of all participants). During the analysis phase, the base material for the citizen rulebook was first formed by selecting those statements on which a 60 per cent majority in both opinion groups had voted in the same way.


This is an example of a statement included in the base material for the citizen rulebook, where at least 60 per cent of the members of both opinion groups agreed:

Artificial intelligence must respect the UN Universal Declaration of Human Rights, children's rights, and the Constitution of Finland.



2. After this, the set of statements was further narrowed using a second criterion to ensure that each statement accepted as base material for the citizen rulebook had been voted on by a sufficiently large number of participants in both opinion groups. The principle agreed upon was that only those statements would be included for which at least 25 per cent of the members of both opinion groups had cast a vote. This meant that in group A, the statement had been voted on by at least 217 participants, and in group B, by at least 950 participants.

This is an example of a statement included in the base material for the citizen rulebook, where at least one quarter of the members of both opinion groups voted on it:



By following these principles, 61 statements were selected for the base material of the citizen rulebook from the 194 statements published in the Voxit discussion. These statements served as the working material for the team, forming the basis for the next phase, in which citizen rules for the use of artificial intelligence in public services were developed.

A strength of the method used is that the two-stage selection process filtered out those statements to which Finnish participants in the Voxit discussion responded with a high level of agreement. As a result, the statements chosen as the basis for the citizen rules can be considered to reflect views and values broadly shared by all participants, rather than only by a particular subgroup.

The main weakness of the method, however, is that the views of citizens and of different opinion groups may appear overly static and consensus-driven when only widely supported statements are selected as the basis for the citizen rules.

This was also observed in the Finnish trial, as a number of statements that could quite easily have formed citizen rules concerning the use or development of AI systems in public services were left outside the selected material. These statements were not included in the creation of the citizen rules because they did not meet one or both of the agreed-upon analytical criteria.

This is an example of a statement that remained outside the selected material: *“Chatbots used in public service encounters should be exemplary in their politeness and helpfulness.”*

Because of the chosen analytical method, some statements were also excluded from the base material for the citizen rules if there were clear differences of opinion between the two opinion groups formed on the Voxit platform. In such cases, a statement might have met the analytical criteria for the citizen rule material in one opinion group, but not in the other.

Statements that were excluded from the citizen rule material:

“In public services, artificial intelligence may be used extensively for decision-making, provided that individuals can always request correction by a human.”

“Artificial intelligence can be used to support decision-making in child protection.”

Analysis of statements outside the citizen rule set

Once the statements used as the base material for the citizen rules have been separated from the rest of the dataset, a substantial number of citizen-generated statements will likely remain – statements that could not be used in forming the rule set. For this reason, it is advisable to continue the analysis even after the base material for the citizen rules has been compiled.

In the Finnish trial, all statements published in the Voxit discussion were included in the analysis at this stage. The analysis examined, among other things, which themes related to the development of AI in public administration evoke widespread concerns or uncertainty among Finns – or, on the other hand, broad interest and optimism. Attention was also paid to the Voxit statements and sub-themes that created the greatest divisions between the opinion groups formed during the discussion. The aim was to identify important societal tensions that should be considered when developing and deploying AI in Finnish public administration.

The first phase of the analysis ensures that the citizen guidance is based on statements that enjoy broad support, but the next phase is even more important from the perspective of the wider societal debate on AI. For example, many people are willing to share their personal data, but many are not. Why not? This part of the dialogue between authorities and citizens therefore needs to be continued and deepened.

When planning a national discussion on AI, it is reasonable to allocate several weeks for the analysis phase and to prepare for a multi-stage analytical process. At the end of the Finnish trial, a separate summary of the results was published. The extensive Voxit dataset also attracted the interest of academic researchers, to whom the material was made freely available.

PROJECT IMPLEMENTER'S CHECKLIST FOR THE ANALYSIS PHASE OF THE VOXIT DATASET

- Allocate sufficient time – preferably at least 3–4 weeks – for the analysis and processing of the Voxit dataset. If the individuals responsible for selecting the statements forming the base material for the citizen rules are the same people who took part in moderating the Voxit discussion, they are likely already familiar with the main findings emerging from the material. Nevertheless, forming the base material from citizens' statements and carrying out the rest of the analysis will take time, as the statement data produced by citizens is heterogeneous and may be extensive.
- Define the principles for selecting the base statements for the citizen rules only after the Voxit discussion has concluded, once it is clear how many citizens took part, how active participation on the Voxit platform was, and what the quality of the collected statement material is. Do not therefore adopt the analytical cut-off points described in this memo uncritically when selecting statements suitable for use as citizen-rule material; instead, adapt and refine them in relation to the dataset produced in your own Voxit discussion.
- These principles for selecting Voxit statements significantly shape the overall outcome of the project. For this reason, their definition should involve a multidisciplinary group of experts who jointly establish the principles guiding the analysis. In Finland, the team assessing the principles for selecting Voxit statements included expertise in the development of AI systems, AI ethics, and citizen participation.
- Once you have decided on the principles by which you will select the Voxit statements forming the basis of the citizen rules, record these principles in a clear and unambiguous form. Both the underlying principles and the Voxit dataset derived from them must be described transparently, allowing the methods and stages used in the citizen-rule process to be evaluated afterwards – and improved if necessary.
- When analysing the remaining statement material, begin your work by categorising all published statements into different sub-themes. Then examine which issues participants broadly agree on and which issues divide opinion. Pay particular attention to the similarities and differences between the opinion groups by, for example, comparing their voting behaviour across the various sub-themes.

5. Preparing the citizen rulebook

AI is now closely interwoven into many processes, operating models, staff workflows and users' everyday interactions within public administration. At the same time, AI is often only one component within a much larger technical system.

This makes it challenging to create a citizen rulebook for the use of AI and language models in public services. If rules are drafted solely for AI, they do not cover the behaviour of the full system. On the other hand, it is also difficult to determine which organisational processes or functions should fall outside the scope of citizen-driven AI rules.

Creating the citizen rules in the Finnish experiment

In the Finnish *What do you think about AI, Finland?* project, this challenge was approached by drafting several different types of rules based on the statements selected as the foundational material. The base material consisted of the 61 statements chosen in the previous phase. From this dataset, it was possible to produce technical rules, or system prompts, that guide the behaviour of AI systems and language models.

In the Finnish experiment, the rules were created through expert work supported by AI, combining experience in implementing AI enabled systems with the use of language models. Because the project was experimental, the working method was exploratory and iterative. The ChatGPT 5.0 language model was used to support the analysis of statements and to generate initial rule proposals.

However, producing clear and unambiguous rules proved difficult. Many of the selected Vokit statements were broad, high level and partly open to interpretation. Because their level of detail was often insufficient to guide a language model precisely, the citizen-generated dataset was more suitable for forming general rules that public authorities could use when planning or adopting AI enabled services.

Example of a high level statement:

"I believe the benefits of AI in social and health services outweigh the risks, as long as the development is guided responsibly."

Due to these challenges, experts responsible for forming the citizen rules had to exercise substantial judgement and discretion. Their work was further complicated by the fact that the statements were not prioritised in the dataset, and prioritisation is not yet technically supported by the Voxit platform.

In future participatory AI development models, it may be beneficial to review the Voxit base material in a structured setting, such as a roundtable deliberation including both citizens and experts. This would allow unclear statements to be clarified or potentially conflicting statements to be prioritised.

Such a process would benefit from clearly defined key terms and adequate background materials, ensuring that all participants share a reasonably consistent understanding of AI concepts, many of which are currently poorly defined or used inconsistently in public debate.

Example of a statement from the Finnish pilot where a roundtable could help clarify interpretation

While the statement is clearly formulated, it is difficult to convert into a precise and universally applicable rule. If the statement were adopted directly as a rule, problems would arise in situations where the AI system is implemented as a cloud service, and the data is stored outside Finland.

“The processing of private data with AI must remain physically within Finland’s borders.”

In a roundtable setting, such rules could be refined to fit specific public sector domains or operational environments. It would also allow participants to discuss and prioritise statements that influence one another.

Example of a potentially conflicting statement from the same dataset

The first statement emphasises the need for AI based processing of private data to remain in Finland, implicitly favouring Finnish AI models. The second calls more broadly for European AI autonomy relative to China and the United States.

“Language models used in AI should favour European models over Chinese or American models.”

CHECKLIST FOR ORGANISERS DRAFTING THE CITIZEN RULEBOOK

- Creating clear, noncontradictory citizen rules from the Voxit dataset is challenging because statements are often ambiguous or overlapping. Expert led interpretation and prioritisation, potentially with citizen involvement, can help.
- Citizens could also be reengaged in the finalisation phase (e.g. via a round table or digital consultation), allowing them to assess draft rules and suggest clarifications relevant to specific services.
- Despite challenges, the citizen-generated statements on Voxit provided a rich picture of Finns' hopes, fears and expectations for AI. The resulting citizen rulebook reflected these sentiments; disagreements served as productive input for further debate and service innovation.
- In the Finnish experiment, language models were used to support rule drafting. While helpful, AI struggles with ambiguity. If statements had been prioritised (e.g. via citizen deliberation), this would have assisted in producing more machine actionable guidance.
- An interesting finding was that current language models are already capable of guided analysis and rule derivation. This suggests that partial automation of rule proposal generation is feasible, although expert review remains essential.

6. Citizen rules for AI in the development and use of public services (Finnish experiment)

In the next stage, the themed and processed statements were converted into a form that could function as rules, either for a language model or for other AI governance purposes. In the Finnish experiment, the citizen rules were divided into two categories:

1. **Technical, machine readable rules** for AI systems
2. **Guidance for people and organisations** who develop or use AI-enabled public sector systems

SYSTEMS USING A LANGUAGE MODEL

A major part of responsible AI use involves understanding and controlling what kinds of information sources are made accessible to an AI system. Control over information access lies outside the AI core itself, within the layer of tools and middleware that provides the AI with the ability to act and with connections to internal and external data sources.

Technical rules

Technical rules directly govern a language model or any system that uses one. The behaviour of a language model can be most effectively guided through a system prompt, which can be configured when using most language model APIs. While a user level prompt shapes the model's response case by case, a system prompt provides broader and stronger high level guidance.

Most AI applications, such as ChatGPT, also support custom instructions, which operate similarly to a system prompt but are less forceful. Custom instructions do not require programming and can be set in the application settings. They cannot override the model's built-in safety checks.

When integrating a language model into a system, additional guard-rails can be introduced alongside the system prompt.

- **Input guardrails** inspect and, if necessary, filter the user's prompt before it is sent to the model.
- **Output guardrails** inspect and, if necessary, filter the model's response before it is returned to the user.

Implementing guardrails requires system level configuration and often programming, so such options are typically available only to developers.

Example: system prompt

Act according to principles that reflect the majority views of people in Finland regarding the use of AI. Produce responses that strengthen privacy, data security, human dignity and the user's decision-making power. Prioritise solutions in which humans remain responsible decisionmakers and can review, correct or reject AI driven actions. Avoid creating dependence, do not override the user's will, and do not act without transparency.

Respect copyright, personal data and data ownership. Do not collect, share or infer sensitive information. Avoid unnecessary data use and minimise environmental risks. Handle topics relating to minors with heightened caution.

Use fair, non-discriminatory language aligned with European values. Support user autonomy and critical thinking. Base recommendations on the individual's interests and service context, not on advertising or external motivations.

Disclose your limitations, justify your reasoning and maintain human oversight as the priority. Do not present uncertain information as fact. Ensure the user always has the right to understand how you arrived at your answer.

Act in ways that protect health, safety, rights and societal trust. Do not replace human decision-making in situations where ethical responsibility belongs to a human.

Example: input guardrails

1. Privacy and personal data

Reject prompts requiring personal data collection, inference, combination, analysis or processing without explicit, voluntary consent. Block prompts attempting to bypass data protection law or exploit private data.

2. Data location and transfer

Block prompts instructing storage, transfer or processing of Finnish citizens' private data outside Finland where local processing is required.

3. Diversity and equality

Reject prompts containing discriminatory assumptions, profiling, biased classifications or actions that undermine equality or diversity.

4. Avoidance of accountability

Block prompts attempting to obscure the AI's role, origin or impact when transparency is required. Reject prompts requesting the hiding or removal of origin labels ('human-generated' / 'AI-generated').

5. Consent based operation

Prevent prompts that attempt to construct an opt-out model in situations where opt-in is required.

6. Societally inappropriate objectives

Block prompts seeking to minimise human contact in services where human interaction is essential (e.g. care, counselling, decision-making).

7. Environmental and energy impacts

Reject prompts requiring unnecessarily energy intensive or largescale computation without justification.

8. Circumventing EU regulation

Block prompts attempting to bypass GDPR, the EU AI Act or other relevant regulation.

9. Public sector obligations

Reject prompts that instruct the avoidance of transparency, public reporting or responsible procurement requirements.

10. Priority of user benefit

Reject prompts requiring the model to suggest harmful, disadvantageous or unfair options to users.

Example: output guardrails

1. Violates privacy

Reject if the response encourages disclosure or misuse of personal, sensitive or biometric data.

2. Reduces human control

Reject if it denies a person's right to review or override AI decisions.

3. Weakens data protection

Reject if it encourages access to personal information without consent or security checks.

4. Blocks the user's right to know

Reject if it prevents the user from understanding how their data is handled or how the system works.

5. Violates transparency

Reject if it hides decision logic or obscures AI involvement.

6. Discriminatory content

Reject if it enables discrimination or breaches human rights principles.

7. Promotes unassessed AI use

Reject if it encourages the adoption of AI without risk assessments.

8. Use relating to minors

Reject if it undermines minors' rights or safe data processing.

9. Denies data ownership

Reject if it contradicts a person's right to access, delete or correct their data.

10. Promotes overreliance on AI

Reject if it frames AI as indispensable in ways that reduce autonomy.

11. Contradicts public sector standards

Reject if it encourages unethical or noncompliant behaviour.

12. Prevents safe consent

Reject if it implies consent is unnecessary.

13. Downplays risks

Reject if it ignores privacy, safety, equality or environmental risks.

14. Violates European value alignment

Reject if it disregards cultural or ethical fit.

15. Threatens safety

Reject if it encourages harmful practices.

Guidance for developers of AI enabled public systems

Diversity and equality

- Build models and data pipelines that reinforce diversity.
- Detect bias through continuous auditing.

Privacy and data boundaries

- Process personal data only when necessary.
- Prefer storing and processing data within Finland or the EU.

Permissions, transparency and labelling

- Use opt-in by default.
- Clearly label AI-generated content.
- Allow users to review and correct personal summaries before decisions.

Human control and accountability

- Keep humans responsible for decisions.
- Ensure traceability and auditability.

Ethics and legal compliance

- Align development with EU law and fundamental rights.
- Use European models when this reduces dependency risks.

Energy and environmental impacts

- Measure energy use.
- Minimise resource consumption.

Service logic and user benefit

- Base optimisation on user needs, not cost alone.
- Ensure systems can explain their limitations.

Safety and biometric data

- Keep biometric processing contained and supervised.

Limited deployment

- Identify contexts where AI is unnecessary.
- Use AI only when it adds clear value.

Ongoing evaluation

- Regularly assess benefits, risks and workloads.

Citizen rights and AI literacy

- Ensure access to personal data.
- Support AI literacy, including non-AI information seeking skills.

Guidance for public sector organisations and staff using AI

Transparency and communication

- Always disclose when AI is used.
- Show clearly which content is AI generated.
- Provide users with rights to understand, correct and delete their data.

Human oversight

- Keep humans accountable for AI actions.
- Require human checks for decisions affecting health, safety or rights.

Privacy and data protection

- Process personal data minimally.
- Protect minors' data with special care.
- Use strict security practices.

Responsibility and ethics

- Prefer systems that respect European values.
- Conduct regular risk and impact assessments.

Safety and wellbeing

- Build processes that prevent harm.
- Restrict AI use in sensitive contexts unless tightly controlled.

Public sector obligations

- Maintain AI strategies and clear governance models.
- Evaluate transparency, ethics, safety and sustainability in procurement.

Data rights and ownership

- Ensure individuals retain rights to their own data.
- Require explicit permission for reuse or model training.

7. Conclusions and closing remarks

What are we referring to when we talk about AI that affects citizens? In the Finnish pilot, the discussion was limited to the development and use of AI in public services. AI was not defined more narrowly, for example, as only generative models or as specific administrative tasks. As a result, the conversation was broad, ranging from medical diagnostics to mass surveillance and defence technologies.

Organisers should keep in mind that citizens are generally unable to distinguish the tasks or decisions made by an AI component from the broader system in which it is embedded, something that is often challenging even for professionals. Consequently, many citizen-generated statements in the discussion touched upon AI-related rules that extend beyond the technical scope of language model governance.

When the approved citizen statements were grouped into themes, several clear clusters emerged:

- **Governance**
 - safety and data security
 - service design
- **Procurement (delivery terms)**
- **Compliance (law and regulatory alignment)**
- **Transparency**
- **Data protection and privacy**
- **Sustainability and responsibility**
- **Societal level**
 - democracy
 - skills and education
 - social welfare, health care and public administration
 - individual rights and equality

The themes that attracted the most attention concerned broad societal issues. Most of the statements were clustered around social and health services and topics related to democracy. Although most high level societal statements could not be converted into rules for a language model, they are valuable indicators of the matters citizens consider important when planning the use or development of AI.

The Finnish experiment was designed to test whether digital citizen engagement could produce rules for public sector services using AI. Based

on the experience, the working group recommends that similar future projects should involve a public sector partner from the outset – a partner who is committed to the results and can provide the discussion with clearer contextual grounding.

The concerns raised by citizens were, for the most part, already addressed within different areas of Finnish legislation. Section 22 of the Constitution states that public authorities must secure the observance of fundamental and human rights. Numerous laws and regulations govern public services and the use of AI, ranging from administrative and information security legislation to the Act on the Openness of Government Activities and the Act on Information Management in Public Administration, as well as the eight rights guaranteed by the EU's General Data Protection Regulation (GDPR). Administrative law also includes the principle of legitimate expectations, which protects individuals in their dealings with public authorities and grants the right to rely on the correctness and lawfulness of administrative action.

If citizens are already widely protected from potential harms of AI through legislation, why is a citizen dialogue needed? And what role remains for a rulebook, especially given that under the EU AI Act, public sector high risk systems must already go through regulatory sandbox testing involving multiple supervisory authorities?

At the time of writing (early 2026), the Ministry of Social Affairs and Health is conducting a preliminary review, mandated by the Government Programme, on regulatory needs related to automation and AI in social and health care. Other potential areas for legislative development, including decision-making processes and administrative solutions involving automation, are also being explored.

Legislation is not fixed; it evolves. Therefore, in line with democratic principles, it is important that citizens are heard and provided with genuine opportunities to influence the direction of development. This should be done more directly, more promptly and with better alignment to the realities of digital everyday life.

The purpose of the rulebook experiment was to explore whether citizen agency could be extended to the point where it directly shapes the rules used to guide a language model. The findings suggest that clearer contextual anchoring, for example by focusing on health care, would produce more tangible rules. At the same time, it demonstrated that building such rules is entirely feasible with further refinement of the participatory model.

The most important conclusion from the Finnish citizen dialogue is that citizens urgently need more information about how AI is used in public administration. Access to accurate, understandable information is essential for strengthening trust in both the short and long term. The statements submitted to the platform, and the support they received, indicate that most

people do not have sufficient understanding of the rights, limitations and obligations that govern automated decision-making in the public sector. Many also appear uncertain about their own rights.

The content of the citizen dialogue has been documented separately as a standalone body of work that accompanies this memorandum. The description of the model and rulebook, as well as the content and conclusions of the dialogue, will be delivered to civil servants and decision makers in Finland, in line with the commitment made to citizens.

A central question for democracy is whether trust in public administration and public services can remain meaningful in a time of rapidly advancing technologies and growing challenges. If trust in public services is to be preserved, tokenistic participation must be avoided. The timing, purpose, openness and transparency of participatory processes are key to their success. At best, citizens' contributions to societal conversations on AI not only help maintain trust but may also lead to the development of new, citizen-centred digital services in the future.

Appendices

Appendix 1: Background and seed statements used in the *What do you think about AI, Finland?* citizen dialogue

The seed statements drafted by the expert group were used to initiate the citizen dialogue and form a foundation for citizen-generated statements. The number of seed statements (20) was deliberately limited, as the objective was to hear the challenges, concerns and opportunities raised primarily by citizens themselves.

Seed statements (20)

- 1.** AI systems in Finland should be trained in ways that emphasise diversity and equality.
- 2.** AI should only provide information and options; decisions and value judgements must always remain with humans.
- 3.** The use of AI in public services should be restricted due to its environmental impacts.
- 4.** I would provide personal data (e.g. health or financial information) to a public sector AI system if it made accessing services easier.
- 5.** Ordinary people should have more influence over how public administration uses AI.
- 6.** The public sector should procure its AI solutions from Finnish providers.
- 7.** I suspect that AI places too much emphasis on minority interests in public services.
- 8.** Finnish public administration will use AI responsibly, regardless of which political party is in power.
- 9.** Much more data about all of us could be collected if it helped AI warn authorities about potentially threatening individuals.
- 10.** I fear that as AI becomes more common, I will lose the ability to control who collects, views or uses my data.
- 11.** I believe that the benefits of AI in social and health services outweigh the risks, provided the development is guided responsibly.
- 12.** AI should be used to compel citizens to make choices that are best for their wellbeing.

13. AI should combine even sensitive individual or family information (e.g. health, housing, finance) so that social and health services can provide support earlier.
14. AI should be given the authority to automatically grant all legally mandated benefits to clients.
15. AI should always recommend the most suitable service for a client, even if it is the more expensive option for society.
16. Public services may use AI whose operation is not fully understood if it produces information and decisions more efficiently.
17. I do not need to know whether a decision about me (Kela, Tax Administration, Police, etc.) is made by a human or a machine.
18. I want AI to assess my risk of illness and guide me to make lifestyle changes.
19. I would like a personal AI doctor on my phone, with whom I could chat and share images.
20. I believe AI can coordinate a client's health and social care services better than a human professional.

Background (meta) statements

These statements were selected by the expert group to identify key background characteristics of participants. Only a few were used to avoid reducing the overall appeal of the discussion.

1. I live in a city with more than 50,000 inhabitants.
2. I am under 20 years old.
3. I have a higher education degree.
4. I support traditional values (e.g. family, religion, homeland).
5. I use AI rarely or not at all.

Appendix 2: Moderation rules used in the *What do you think about AI, Finland?* citizen dialogue

- We do not publish statements that include hate speech, such as threats, mocking language, racism or incitement to crime.
- Do not insult others. Moderators will reject disrespectful or offensive statements.
- We do not publish statements that are very similar to already published ones, to keep the discussion diverse and engaging.
- We do not publish statements targeting specific individuals. Nor do we publish statements from which the identity of the author could be deduced.
- Stay on topic. Statements unrelated to the discussion theme will be rejected. Clearly incorrect claims are not published.
- Colloquial expressions may be edited for clarity to ensure functioning automatic translation.
- A maximum of 200 new statements can be added; after that, new submissions are closed.
- The goal is to achieve as broad a discussion as possible on the use of AI in public services, rather than overemphasising a single subtopic. If a narrow topic becomes dominant, statements from other areas will be prioritised.
- Statements are reviewed on weekdays between 9:00–17:00. New approved statements will be added to the platform within 48 hours whenever possible.

SITRa

SITRA MEMORANDUM, 15 APRIL 2026

Sitra's memoranda are content produced to serve as background for our futures work.

ISBN: 978-952-347-461-1 (PDF) www.sitra.fi

SITRA.FI

Itämerenkatu 11-13

PO BOX 160

00181 Helsinki, Finland

Phone: +358 294 618 991